Informational Braess' Paradox: The Effect of Information on Traffic Congestion

Daron Acemoglu* Ali Makhdoumi† Azarakhsh Malekian‡
Asuman Ozdaglar§

January 12, 2016

arXiv:1601.02039v1 [cs.GT] 8 Jan 2016

Abstract

To systematically study the implications of additional information about routes provided to certain users (e.g., via GPS-based route guidance systems), we introduce a new class of congestion games in which users have differing information sets about the available edges and can only use routes consisting of edges in their information set. After defining the notion of Information Constrained Wardrop Equilibrium (ICWE) for this class of congestion games and studying its basic properties, we turn to our main focus: whether additional information can be harmful (in the sense of generating greater equilibrium costs/delays). We formulate this question in the form of Informational Braess' Paradox (IBP), which extends the classic Braess' Paradox in traffic equilibria, and asks whether users receiving additional information can become worse off. We provide a comprehensive answer to this question showing that in any network in the series of linearly independent (SLI) class, which is a strict subset of series-parallel network, IBP cannot occur, and in any network that is not in the SLI class, there exists a configuration of edge-specific cost functions for which IBP will occur. In the process, we establish several properties of the SLI class of networks, which are comprised of linearly independent networks joined together. These properties include the characterization of the complement of the SLI class in terms of embedding a specific set of subgraphs, and also show that whether a graph is SLI can be determined in linear time. We further prove that the worst-case inefficiency performance of ICWE is no worse than the standard Wardrop Equilibrium with one type of users.

1 Introduction

The advent of GPS-based route guidance systems, such as Waze or Google maps, promises a better traffic experience to its users, as it can inform them about routes that they were not aware of or help them choose dynamically between routes depending on recent levels of congestion. Though other drivers might plausibly suffer increased congestion as

*Department of Economics, Massachusetts Institute of Technology, Cambridge, MA, 02139 . daron@mit.edu

†Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, 02139 makhdoum@mit.edu

‡Rotman School of Management, University of Toronto, Toronto, ON, M5S 3E6 azarakhsh.malekian@rotman.utoronto.ca

§Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, 02139 asuman@mit.edu

the routes they were using become more congested due to this reallocation of traffic, or certain residents may experience elevated noise levels in their side streets, it is generally presumed that the users of these systems (and perhaps society as a whole) will benefit. In this paper, we present a framework for systematically analyzing how changes in the information sets of users in a traffic network (e.g., due to route guidance systems) impact the traffic equilibrium, and show the conditions under which even those with access to additional information may suffer greater congestion.

Our formal model is a version of the well-known congestion games, augmented with multiple types of users (drivers), each with a different information set about the available edges in the network. These different information sets represent the differing knowledge of drivers about the road network, which may result from their past experiences, from inputs from the social network, or from the different route guidance systems they might rely on. A user can only utilize a route (path between origin and destination) consisting of edges belonging to her information set. Generalizing the classic notion of *Wardrop equilibrium* (Wardrop (1952), and Beckmann et al. (1956)), where each (non-atomic) user takes the level of congestion on all edges as given and chooses her optimal route, we define the notion of *Information Constrained Wardrop Equilibrium (ICWE)*, which also imposes the same equilibrium condition as Wardrop equilibrium, but only for routes that are contained in the information set of each class of user.

After establishing the existence and essential uniqueness of ICWE and characterizing its main properties for networks with a single origin-destination pair (an assumption we impose for simplicity), we turn to our key question of whether expanding the information sets of some groups of users can make them worse off — in the sense of increasing the level of congestion they suffer in equilibrium. For this purpose, we define the notion of *Informational Braess' Paradox (IBP)*, designating the possibility that users with expanded information sets experience greater equilibrium congestion. We then provide a tight characterization of when IBP is and is not possible in a traffic network.

Our main result is that IBP does not occur if and only if the network/graph is in the *series of linearly independent (SLI)* class. More specifically, this result means that in an SLI network, IBP can never occur, ensuring that users with expanded information sets will benefit from their additional information; and conversely that if the network is not SLI, then there exists a configuration of latency/cost functions for edges for which IBP will occur. To understand this result, let us consider what the relevant class of networks comprises. The set of SLI graphs is a subset of *series-parallel* graphs, which are those consisting of parallel paths that may comprise several edges and may intersect, but each intersecting edge is traversed only in one direction. An SLI graph is obtained by joining together a collection of *linearly independent (LI)* graphs, and an LI graph has the crucial property that for each route there exists at least one edge which only belongs to that route — hence giving the graph a linearly independent representation when cast in the form of vectors in the edge space. This in particular implies that in the traffic equilibrium with heterogeneous users over a LI or SLI graph, if we reduce the total flow, then there will always exist one route for which the total cost/latency will strictly go down (this property does not hold in general graphs and not even in series-parallel graphs, and SLI graphs makes up the largest class for which this property is true). A crucial step in our mathematical argument is that when certain users gain additional information, they will change their routing, say redirecting it to some subgraph A of the original graph from some other subgraph B (and since the original graph is SLI, both A and B are also SLI). All else equal, this will increase flows and costs in A and reduce it in B. Without

any adjustment from other users, this cannot increase the costs faced by the users with additional information, since they decrease flows in B must a fortiori reduce flows and costs in some routes in B by the SLI property, and the users in question have access to the routes in B, so their costs must weakly decrease. What happens when other users adjust? Since flows and thus costs in A increase, there will be a reallocation away from A, but again by the SLI property, this implies that costs in A must weakly decrease, and once more because the users in question now have access to routes in A, their costs cannot increase. This proves the "if" part of our main result. The "only if" part is proved by establishing that every non-SLI graph embeds one of the collection of subgraphs, and we can show constructively that each one of these subgraphs is sufficient to generate IBP.

We should also note that, since SLI is a very restrictive class of networks, and few real-world networks would fall into this class, we take this characterization to imply that IBP is very difficult to rule out, and thus the new and highly-anticipated route guidance technologies may make traffic problems worse.

Since the class of SLI networks plays a central role in our analysis and in determining whether IBP occurs, a natural question is whether identifying SLI networks is straightforward. We answer this question by showing that whether a given network is SLI or not can be determined in linear time.

If, rather than considering a general change of information sets, we specialize the problem so that the change in the available information brings all users to complete information about the available set of routes, then we can show that an IBP is possible if and only if the network/graph is *series-parallel*. It is intuitive that this class of graphs is less restrictive than SLI, since we are now considering a very specific change information sets (thus making it less likely that an IBP can occur).

Our notion of IBP relates closely to the classic Braess' Paradox, which considers when all users' costs will increase when an additional edge is added to a network (Braess (1968); Murchland (1970); Arnott and Small (1994)). BP can be seen as a special case of IBP when all users have the same information set to start with and an edge is added to their information set. Our characterization of ICWE and IBP clarifies that our notion is different and, at least mathematically, more general. This can be seen readily from a comparison of our results to the most closely related papers to ours in the literature, Milchtaich (2005, 2006). Milchtaich (2006)'s characterization implies that BP can be ruled out in series-parallel networks. Since IBP is a generalization of BP, it should occur in a wider class of networks, and this is indeed what our result shows in view of the fact that SLI is a strict subset of series-parallel graphs. This result also indicates that IBP is a considerably more pervasive phenomenon than BP. The mathematical argument for our key theorem is also different from Milchtaich (2006) due to the key difficulty relative to BP that not all users have access to the same set of edges, and thus changes in traffic that benefit some groups of users might naturally harm others by increasing the congestion on the paths that they were previously utilizing.

Issues related to Braess' Paradox arise not only in the context of models of traffic, but in various models of communication, pricing and choice over congested goods, and electrical circuits (e.g., Orda et al. (1993), Korilis et al. (1997), Kelly et al. (1998), and Low and Lapsley (1999) for communication networks; the classic works by Pigou (1920) and Samuelson (1952) as well as more recently, Johari and Tsitsiklis (2003), Acemoglu and Ozdaglar (2007); Acemoglu et al. (2007); Ashlagi et al. (2009) and Perakis (2004) for related economic problems; Frank (1981), Cohen and Horowitz (1991), and Cohen and Jeffries (1997) for mechanical systems electrical circuits; and Rosenthal (1973) and Vetta (2002) for gen-

eral game-theoretic approaches). This observation also implies that the results we present here are relevant beyond traffic networks, in fact to any resource allocation problem over a network subject to congestion considerations.[1]

Because our analysis also presents the "price of anarchy," i.e., bounds on the overall level of inefficiency that can occur in an ICWE, our paper is related to previous work on price of anarchy in congestion and related games, e.g., Koutsoupias and Papadimitriou (1999), Roughgarden and Tardos (2002), Correa et al. (2004, 2005), , and Friedman (2004), as well as more generally to the analysis of equilibrium and inefficiency in various variants of this class of games, including Steinberg and Zangwill (1983), Dafermos and Nagurney (1984), Milchtaich (2004a,b), Schoenmakers (1995); Eppstein (1992); Mavronicolas et al. (2007); Feldman and Friedler; Nisan et al. (2007); Rogers et al. (2015); Murchland (1970); Arnott and Small (1994); Lin et al. (2004); Feldman and Friedler (2015); Meir and Parkes (2015); Nikolova and Stier-Moses (2014), Sanghavi and Hajek (2004), Maheswaran and Başar (2003), Yang and Hajek (2005), Anshelevich et al. (2008), , Patriksson (1994), and Meir and Parkes (2014). Here, our result is that the presence of users with different information sets does not change the worst-case inefficiency traffic equilibrium as characterized, for example, in Roughgarden and Tardos (2002).

The rest of the paper is organized as follows. In the next section, we introduce our model of traffic equilibrium with users that are heterogeneous in terms of the information about routes/edges they have access to, and then define the notion of Information Constrained Wardrop Equilibrium for this set up. In Section 3 we prove the existence and essential uniqueness of Information Constrained Wardrop Equilibrium. Before moving to our main focus, in Section 4 we review some graph-theoretic notions about series-parallel and linearly independent graphs, and then introduce the class of series of linearly independent graphs and prove some basic properties of this class of graphs, which are then used in the rest of our analysis. Section 5 defines our notion of Informational Braess' Paradox. Section 6 contains our main result, showing that Informational Braess' Paradox occurs "if and only if" the network is not in the class of series of linearly independent graphs. Section 7 characterize is the worst-case inefficiency of Information Constrained Wardrop Equilibrium, and finally, Section 8 concludes.

2 Model

We first describe the environment and then introduce our notion of Information Constrained Wardrop Equilibrium.

2.1 Environment

We consider an undirected multigraph consisting of a finite vertex set V and a finite edge set $\mathcal{E}$. Each edge $e \in \mathcal{E}$ joins two distinct vertices, u and v, which are referred to as the end vertices of e. Thus, loops are not allowed, but more than one edge can join two vertices (that is the reason that we use the term multigraph). An edge e and a vertex v are said

[1] As pointed out in Newell (1980) and Sheffi (1985), the Braess' paradox and related inefficiencies are a clear and present challenge to traffic engineers, who often try to restrict travel choices to improve congestion (e.g., via systems such as ramp metering on freeway entrances (Sheffi (1985))). We should also note that Braess' paradox has been observed beyond traffic networks, for example, in computer and telecommunication networks, electrical circuits, and mechanical systems (Frank (1981), Cohen and Horowitz (1991), Cohen and Jeffries (1997)).

to be *incident* with each other if v is an end vertex of e. A path of length n ($n \geq 0$) is an alternating sequence p of vertices and edges, $v_0e_1v_1 \dots v_{n-1}e_nv_n$, beginning and ending with vertices, in which each edge is incident with the two vertices immediately preceding and following it and all the vertices (and necessarily all the edges) are distinct. We may also show a path by a series of edges $e_1, \dots, e_n$ knowing that the ending vertex of e_i is the same as the beginning vertex of e_{i+1}, for $i = 1, \dots, n-1$. The first and last vertices, v_0 and v_n, are called the initial and terminal vertices in p, respectively. If q is a path of the form $v_ne_{n+1}v_{n+1} \dots e_mv_m$, with the initial vertex same as the terminal vertex of p but all the other vertices and edges are not in p, then $v_0e_1v_1 \dots e_nv_ne_{n+1}v_{n+1} \dots e_mv_m$ is also a path, denoted by $p + q$.

Throughout, we focus on a two-terminal network (network, for short) $G = (V, \mathcal{E})$, which is an undirected multigraph together with an ordered pair of distinct vertices, an origin O and a destination D.[2] When this causes no confusion, we will refer to G either as a graph or as a network. Any path r with these initial and terminal vertices will be called a route. The set of all routes in a network is denoted by $\mathcal{R}$, i.e.,

$$\mathcal{R} = \{r \ : \ r \text{ is a path with initial vertex } O \text{ and terminal vertex } D\}.$$

We suppose that there are K types of users (we use the terms users and players interchangeably) denoted by $[K] = \{1, \dots, K\}$. Player $i \in [K]$ has a prescribed amount of traffic load $s_i \geq 0$. We also denote the set of all traffic loads by $s_{1:K} = (s_1, \dots, s_K)$. For each player i, we use $\mathcal{E}_i \subseteq \mathcal{E}$ to denote the set of edges that are known to users in type i. The set $\mathcal{E}_i$, $i = 1, \dots, K$, represent agent i's information regarding available edges in the network. Because a user can only utilize routes consisting of edges in her information set, the set of routes available to a user of type i is

$$\mathcal{R}_i = \{r \in \mathcal{R} \text{ such that all edges of } r \text{ belong to } \mathcal{E}_i\}.$$

We refer to $\mathcal{E}_i$ and $\mathcal{R}_i$ as the edges and routes in type i's information set. We use $\mathcal{E}_{1:K} = (\mathcal{E}_1, \dots, \mathcal{E}_K)$ to denote information sets of all types.
We use $f^{(i)} = (f_r^{(i)} \mid r \in \mathcal{R}_i)$ to denote the flow vector of type i, where for all $r \in \mathcal{R}_i$, $f_r^{(i)} \geq 0$ represents amount (flow) of traffic that type i sends on route r. We use $f^{(1:K)} = (f^{(1)}, \dots, f^{(K)})$ to denote the flow vector of all types. Each edge of the network has a cost (latency) function $c_e : \mathbb{R}^+ \to \mathbb{R}^+$ that is continuous, nonnegative, and nondecreasing. We denote the set of all cost functions by

$$\mathbf{c} = \{c_e \ : \ e \in \mathcal{E}\}.$$

For instance, if all the cost functions are affine functions, then for any $e \in \mathcal{E}$, we would have $c_e(x) = a_ex + b_e$, for some $a_e, b_e \in \mathbb{R}^+$.

Definition 1 (Traffic Network with Multiple Information Types) We refer to $(G, \mathcal{E}_{1:K}, s_{1:K}, \mathbf{c})$ as a traffic network with multiple information types. A feasible flow is a flow vector $f^{(1:K)} = (f^{(1)}, \dots, f^{(K)})$ such that for all i, $f^{(i)}$ is an assignment of the traffic load s_i into the routes available to type i, i.e., $f^{(i)} : \mathcal{R}_i \to \mathbb{R}^+$ is a mapping that satisfies

$$\sum_{r \in \mathcal{R}} f_r^{(i)} = s_i,$$

[2] We assume that each vertex and each edge belongs to at least one path between the initial vertex O and the terminal vertex D. This assumption is without loss of generality because the vertices and edges that do not belong to any path from O to D are irrelevant for the purpose of sending traffic from O to D.

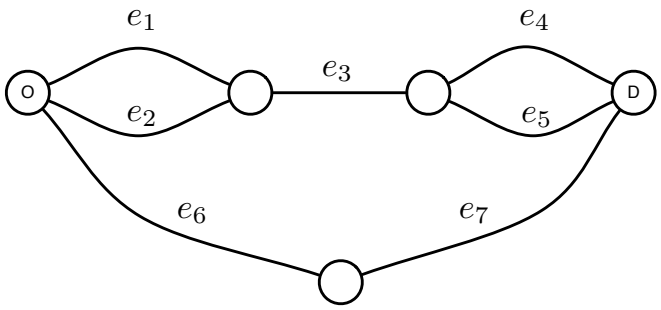

Figure 1: Example of a network with costs given by $c_{e_1}(x) = c_{e_4}(x) = c_{e_6}(x) = x$ and $c_{e_2}(x) = c_{e_5}(x) = c_{e_7}(x) = 1 + ax$ and $c_{e_3} = ax$.

for all $i = 1, \ldots, K$. We denote the total flow on each route r by f_r, i.e.,

$$f_r = \sum_{i=1}^{K} f_r^{(i)}.$$

2.2 Information Constrained Wardrop Equilibrium

The cost of a route with respect to a flow $(f^{(1)}, \ldots, f^{(K)})$ is the sum of the cost of the edges that belong to the route, i.e.,

$$c_r(f^{(1:K)}) = \sum_{e \in r} c_e(f_e),$$

where f_e denotes the amount of traffic that passes through edge e, i.e.,

$$f_e = \sum_{r \in \mathcal{R} \,:\, e \in r} f_r, \; i = 1, \ldots, K.$$

We assume flows get allocated at equilibrium according to a "constrained" version of Wardrop's principle: flows of each type of users are routed along routes in their information set with minimal (and hence equal) cost. We next formalize this equilibrium notion.

Definition 2 (Information Constrained Wardrop Equilibrium (ICWE)) A feasible flow $f^{(1:K)} = (f^{(1)}, \ldots, f^{(K)})$ is an Information Constrained Wardrop Equilibrium (ICWE) if for every pair $r, \tilde{r} \in \mathcal{R}_i$ with $f_r^{(i)} > 0$, we have

$$c_r(f^{(1:K)}) \leq c_{\tilde{r}}(f^{(1:K)}). \tag{2.1}$$

This implies that all routes of type i with positive flow have the same cost, which is smaller or equal to the cost of any other route in $\mathcal{R}_i$. The equilibrium cost of type i, denoted by $c^{(i)}$, is then given by the cost of any route in $\mathcal{R}_i$ with positive flow from type i. Note that Wardrop Equilibrium (WE) is a special case of this definition for a traffic network with a single information type, i.e., $K = 1$.

We next provide an example that illustrates this definition and how it differs from the classic Wardrop Equilibrium.

Example 1 Consider the network given in Figure 1 with $s_1 = s$ and $s_2 = 1 - s$. There are 5 different routes from origin to destination, which we denote by $r_1 = e_1e_3e_4$, $r_2 = e_1e_3e_5$, $r_3 = e_2e_3e_4$, $r_4 = e_2e_3e_5$, and $r_5 = e_6e_7$. We let $\mathcal{E}_1 = \mathcal{E}$, which results in $\mathcal{R}_1 = \{r_1, r_2, r_3, r_4, r_5\}$ and $\mathcal{E}_2 = \{e_6, e_7\}$, which results in $\mathcal{R}_2 = \{r_5\}$.

- If $s \leq \frac{2+a}{3+2a}$, ICWE is $f^{(2)}_{r_5} = 1 - s$ and $f^{(1)}_{r_1} = s$. The equilibrium cost of type 1 is $c^{(1)} = c_{r_1}(f^{(1:2)}) = s + as$. The equilibrium cost of type 2 is $c^{(2)} = c_{r_5}(f) = (1 - s) + 1 + a(1 - s)$. This shows that the equilibrium cost of type 1 and type 2 users need not be the same.
- If $s > \frac{2+a}{3+2a}$, ICWE is $f^{(1)}_{r_1} = \frac{2+a}{3+2a}$, $f^{(1)}_{r_5} = s - \frac{2+a}{3+2a} > 0$ and $f^{(2)}_{r_5} = 1 - s$, which give $c^{(1)} = c^{(2)} = \frac{(2+a)^2}{3+2a}$. This illustrates that when they use a common path, the cost of different types of players are the same.

3 Existence of Information Constrained Wardrop Equilibrium

In this section, we show that given a traffic network with multiple information types $(G, \mathcal{E}_{1:K}, s_{1:K}, \mathbf{c})$ a ICWE always exists and the equilibrium cost of each type of player is uniquely defined, i.e., equilibrium cost for each type is the same for all equilibria. Our proof for existence and "essential" uniqueness of ICWE relies on the following characterization, which is a slight extension of the well-known optimization characterization of Wardrop Equilibrium (see Beckmann et al. (1956); Smith (1979))

Proposition 1 *A feasible flow $f^{(1:K)}$ is a ICWE if and only if it is a solution of the following optimization problem*

$$\begin{aligned} &\min \sum_{e \in E} \int_0^{f_e} c_e(z)dz \\ &f_e = \sum_{i=1}^{K} \sum_{\{r \in \mathcal{R}_i \mid e \in r\}} f^{(i)}_r, \\ &\sum_{r \in \mathcal{R}_i} f^{(i)}_r = s_i, \text{ and } f^{(i)}_r \geq 0 \text{ for all } r \in \mathcal{R}_i. \end{aligned} \tag{3.1}$$

We call $\sum_{e \in \mathcal{E}} \int_0^{f_e} c_e(z)dz$ the potential function and denote it by Φ.

Proof Since the functions c_e for any $e \in \mathcal{E}$ are nondecreasing, the function $\int_0^{f_e} c_e(z)dz$ for each e as a function of $f^{(i)}_r$ is convex and continuously differentiable. We form the Lagrangian function as

$$\mathcal{L} = \sum_{e \in \mathcal{E}} \int_0^{f_e} c_e(z)dz - \sum_{i=1}^{K} \lambda_i \left(\sum_{r \in \mathcal{R}_i} f^{(i)}_r - s_i \right).$$

The first order optimality condition yields

$$\frac{\partial \mathcal{L}}{\partial f^{(i)}_r} = \sum_{e \in E} \frac{\partial f_e}{\partial f^{(i)}_r} c_e(f_e) = \sum_{e \in E \,:\, e \in r} c_e(f_e) = \lambda_i,$$

when $f_r^{(i)} > 0$ and

$$\frac{\partial \mathcal{L}}{\partial f_r^{(i)}} = \sum_{e \in E} \frac{\partial f_e}{\partial f_r^{(i)}} c_e(f_e) = \sum_{e \in E \,:\, e \in r} c_e(f_e) \geq \lambda_i,$$

when $f_r^{(i)} = 0$. These two relations along with $\sum_{e \in E \,:\, e \in r} c_e(f_e) = c_r(f)$, yields (2.1), which is the definition of ICWE. ■

Using the characterization of ICWE as the minimizer of a potential function, we now prove existence and essential uniqueness.

Theorem 1 (Existence and Uniqueness of ICWE) *Let $(G, \mathcal{E}_{1:K}, s_{1:K}, \mathbf{c})$ be a traffic network with multiple information types.*

1. *There exists a ICWE $f^{(1:K)} = (f^{(1)}, \ldots, f^{(K)})$.*
2. *The ICWE is essentially unique in the sense that if $f^{(1:K)}$ and $\tilde{f}^{(1:K)}$ are two constrained Wardrop equilibriums, then $c_e(f_e) = c_e(\tilde{f}_e)$ for every edge $e \in E$.*

Proof The set of feasible flows $(f^{(1)}, \ldots, f^{(K)})$ is a compact subset of $K|\mathcal{R}|$-dimensional Euclidean space. Since edge cost functions are continuous, the potential function is also continuous. Weierstrass extreme value theorem establishes that optimization problem (3.1) which by Proposition 1 is a ICWE, attains its minimum.
We next, show that in two different equilibria $(f^{(1)}, \ldots, f^{(K)})$ and $(\tilde{f}^{(1)}, \ldots, \tilde{f}^{(K)})$, the equilibrium cost for each type is the same. By Proposition 1, both $(f^{(1)}, \ldots, f^{(K)})$ and $(\tilde{f}^{(1)}, \ldots, \tilde{f}^{(K)})$ are optimal solutions of (3.1). Since Φ is a convex function, we have that

$$\begin{aligned} &\Phi(\alpha(f^{(1)}, \ldots, f^{(K)}) + (1-\alpha)(\tilde{f}^{(1)}, \ldots, \tilde{f}^{(K)})) \\ &\leq \alpha \Phi(f^{(1)}, \ldots, f^{(K)}) + (1-\alpha)\Phi(\tilde{f}^{(1)}, \ldots, \tilde{f}^{(K)}), \end{aligned}$$

for any $\alpha \in [0, 1]$. Since $\Phi(f^{(1)}, \ldots, f^{(K)})$ and $\Phi(\tilde{f}^{(1)}, \ldots, \tilde{f}^{(K)})$ are both equal to optimal value of (3.1), and for each e, the function $\int_0^{f_e} c_e(z)dz$ is convex (its derivative with respect to f_e is $c_e(f_e)$ which is non-decreasing), the functions $\int_0^{f_e} c_e(z)dz$ for any $e \in \mathcal{E}$ must be linear between values of f_e and $\tilde{f}_e$. This shows that all cost functions c_e are constant between f_e and $\tilde{f}_e$. ■

Remark 1 The essential uniqueness of equilibrium is not a general property of congestion games with multiple types. In particular, when different classes of users have different cost functions for the same edge, Wardrop Equilibrium may involve different cost levels for different users. For example, Milchtaich (2005); Mavronicolas et al. (2007) provide sufficient and necessary topological conditions on the network under which equilibrium cost is uniquely defined for all users. In our model, essential uniqueness obtains because the cost of an edge e for type i users is c_e if $e \in \mathcal{E}_i$ and ∞ otherwise.

Remark 2 Theorem 1 assumes that the cost functions are only non-decreasing. If we strengthen this assumption to strictly increasing cost functions, then the essential uniqueness result can be strengthened as well. In particular, in this case, the total flow on any edge for all equilibria would be the same. In particular, if $(f^{(1)}, \ldots, f^{(K)})$ and $(\tilde{f}^{(1)}, \ldots, \tilde{f}^{(K)})$ are two equilibria, then for any edge $e \in \mathcal{E}$, we have $f_e = \tilde{f}_e$, where $f_e = \sum_{r \in \mathcal{R} \,:\, e \in r} \sum_{i=1}^{K} f_r^{(i)}$ and $\tilde{f}_e = \sum_{r \in \mathcal{R} \,:\, e \in r} \sum_{i=1}^{K} \tilde{f}_r^{(i)}$.

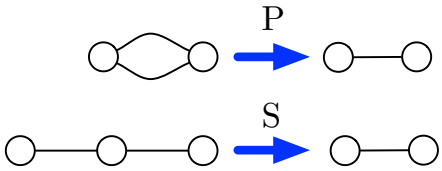

Figure 2: Two operations that turns a series-parallel network into a single edge.

4 Some Graph-Theoretic Notions

In this section, we present several classes of graphs which we use to study in our characterization of IBP. In preparation for our main results, we also present various equivalent characterizations of these graphs and delineate the relations among them. We first define the class of series-parallel networks.

Definition 3 (Series-Parallel Network (SP)) A (two-terminal) network is called series-parallel if two routes never pass through any edge in opposite directions.

As an example, the networks given in Figure 4a and Figure 4b are series-parallel networks, while the network given in Figure 4c is not series-parallel. The reason is that two routes e_1, e_5, e_4 and e_2, e_5, e_3 pass through the edge e_5 in opposite directions. Series-parallel networks can be recursively characterized as was shown by Riordan and Shannon (1942).

For future reference, we also note that identifying series-parallel networks is one of the classical problems in the design of algorithms, and it is well known Valdes et al. (1979) that this can be performed in linear time , i.e., in $O(|\mathcal{E}| + |V|)$. Moreover, Valdes et al. (1979) show the following tree decomposition of a series-parallel network in linear time.

Proposition 2 (Valdes et al. (1979)) *(a) A network is series-parallel if following the procedure shown in Figure 2 in any order, turns the network into a single edge connecting origin to destination. Moreover, this can be done in linear time.*

(b) The binary tree decomposition (shown in Figure 3) obtained by the procedure shown in Figure 2 can be obtained in linear time $O(|\mathcal{E}| + |V|)$.

Proposition 3 (Riordan and Shannon (1942)) *A two terminal network is series-parallel if and only if the network*

(i) comprises a single edge between O and D.

(ii) is created by connecting two series-parallel networks in series.[3]*.*

(iii) is created by connecting two series-parallel networks in parallel.[4]

An important subclass of series-parallel graphs are *linearly independent* graphs.

[3] That is, by joining the destination of one series-parallel network with the origin of the other one

[4] That is, by joining the origins and destinations of two series-parallel graphs.

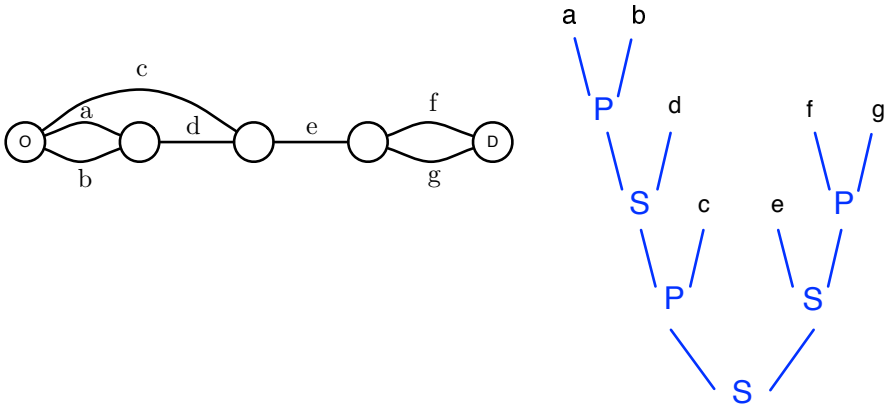

Figure 3: Binary tree decomposition of a series-parallel network.

Definition 4 (Linearly Independent Network (LI)) A (two terminal) network is called linearly independent (LI) if each route has at least one edge that does not belong to any other route.

Remark 3 Linearly independent here refers to the fact that the network G has "linearly independent" routes, meaning that its set of routes constitute a linearly independent set of vectors in the edge space (see Diestel (2000); Milchtaich (2006)). More precisely, for any $r \in \mathcal{R}$, let $v_r \in \mathbb{F}_2^{|\mathcal{E}|}$ be defined as

$$v_r = (v_r^1, \ldots, v_r^{|\mathcal{E}|}), \text{ where } v_r^i = 1 \text{ if } e_i \in r \text{ and } 0 \text{ otherwise.}$$

If the network G is LI, the set of vectors $\{v_r \ : \ r \in \mathcal{R}\}$ is linearly independent.

Similar to series-parallel graphs, LI graphs can also be characterized recursively, as was shown in Holzman and Law-yone (2003).

Proposition 4 (Holzman and Law-yone (2003)) *A two terminal network has linearly independent routes if and only if one of the following conditions holds:*

(i) it is comprised of a single edge between O and D.

(ii) it is the result of connecting two in parallel.

(iii) it is the result of connecting one LI graph with a single edge, i.e., extending either origin or destination of the LI graph.

As it is clear from Propositions 3 and 4 the class of linearly independent networks is a subset of the class of series-parallel networks. The following proposition provides an alternate characterization of linearly independent and series-parallel networks that we will use in the proofs. We need the notion of graph embedding to present these characterizations.

Definition 5 (Embedding) A graph H is embedded in the graph G if we can start from H and create G with the following procedure, which is called embedding.

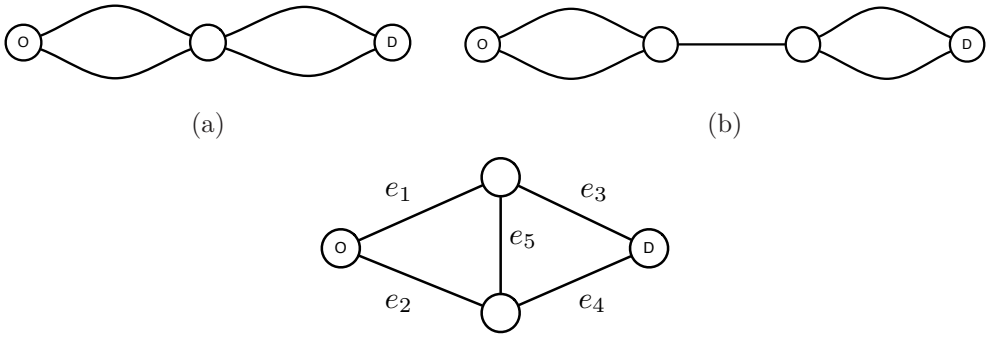

Figure 4: Basic networks that are not embedded in series-parallel networks and linearly independent networks: (c) is not embedded in series-parallel networks. (a), (b), and (c) are not embedded in linearly independent networks.

(i) Divide an existing edge and put a new vertex between the two endpoints of that edge.

(ii) Add an edge between two nodes.

(iii) Extend origin or destination by one edge.

Proposition 5 *We have the following characterizations:*

(a) [Milchtaich (2006)] A network G is linearly independent if and only if none of the networks given in Figure 4 are embedded in it. Furthermore, for every pair of routes r and r' and every vertex v common to both routes, either the section r_{ov} (which consists of v and all the vertices and edges preceding it in r) is equal to r'_{ov}, or r_{vd} is equal to r'_{vd}.

(b) [Duffin (1965)] A network G is series-parallel if and only if the network given in Figure 4c is not embedded in it.

This proposition shows that series-parallel networks are essentially those that do not embed the well-known Wheatstone network shown in Figure 4c. LI networks, in addition, also exclude series-parallel networks that have routes that "touch" as indicated in Figure 4a and Figure 4b.

We now introduce a new class of graphs, *series of linearly independent graphs* (SLI), which will use in our main result below.

Definition 6 (SLI) A (two-terminal) network is called series of linearly independent if it consists of series of networks that are linearly independent. In other words, network G is SLI if

(i) G is a linearly independent network.

(ii) G is the result of connecting two SLI networks in series.

We next provide a new characterization of SLI graphs, extending the characterization for SP and LI graphs presented in Proposition 5.

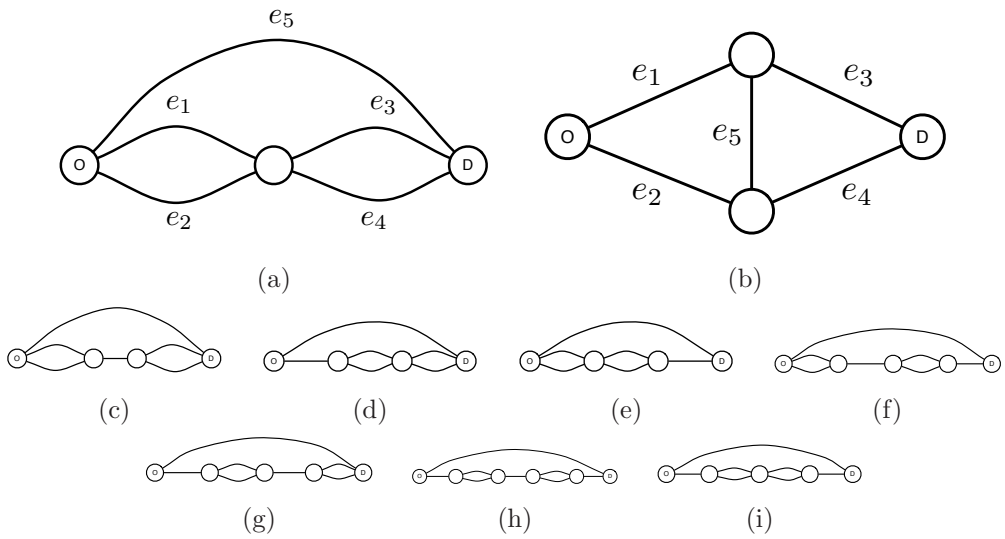

Figure 5: Basic networks that are not embedded in SLI networks.

Theorem 2 (Characterization of SLI) *A network G belongs to the class SLI if and only if none of networks shown in Figure 5 is embedded in it.*

Proof We first show that if a network G belong to the class SLI, then none of the networks shown in Figure 5 is embedded in it. First note that since all networks in the class SLI are series-parallel, using part (b) of Proposition 5 implies that the Wheatstone network shown in Figure 5b is not embedded in it. A SLI network consist of several LI networks that are attached in series. We call each linearly independent subnetwork of G an LI part. We first show that the network shown in Figure 5 can only be embedded in one LI part of G. Let G_1 and G_2 be two LI networks that are attached in series where the resulting network from this attachment is H. Also, let the node c be the attaching point of these two networks. We will show that the network shown in Figure 5a can not be embedded in H (a similar argument shows none of the other networks shown in Figure 5 is embedded in it). In order to reach to a contradiction, we suppose the contrary, i.e., H is obtained from the network shown in Figure 5a by applying the embedding process described in Definition 5. We define the corresponding routes to e_5, e_1e_4, and e_2e_3 in H by r_3, r_1 and r_2. Formally, we start from $r_3 = e_5$, $r_1 = e_1e_4$, and $r_2 = e_2e_3$ and in the embedding procedure whenever we divide an edge on r_i $(i = 1, 2, 3)$ we will update r_i by adding that edge and whenever we extend origin or destination we will add the new edge to r_i. Given this definition, in the network H we have three routes r_3, r_1, and r_2, where r_1 and r_2 have a common node and they do not have any common node with r_3 (except O and D) with r_3. This is a contradiction as all routes in H have node c in common. Therefore, if one of the networks shown in Figure 5 is embedded in G it must be embedded in one of the LI parts, which is a contradiction by using part (a) of Proposition 5, completing the proof of the first part.

We next show that if none of networks shown in Figure 5 is embedded in G, then G belong to the class SLI. Proposition 5(b) implies that since Figure 5b is not embedded in the network G, the network is series-parallel. We next show that given a series-parallel network G, if G is not SLI then we can find an embedding of one of Figure 5a, 5c, 5d, 5e, 5f, 5g, 5h, or 5i in it. The proof is by induction on the number of edges of G. Following

Proposition 3, consider the last building step of the network G. If the last step, is attaching two networks G_1 and G_2 in series, then assuming that G is not SLI, we conclude that either G_1 or G_2 (or both) is not SLI. Therefore, by induction assertion, we can find an embedding of one of the Figure 5a,..., Figure 5i in either G_1 or G_2, completing the proof for the series attachment. Otherwise, the last building step of network G is by attaching two networks G_1 and G_2 in parallel. In this case, first note that since LI $\subset$ SLI, G being not SLI implies that G is not LI. If G is not LI, then either G_1 or G_2 is not LI, otherwise the parallel attachment of them would be LI (using Proposition 4). Let the network that is not LI be G_1. Therefore, part (a) of Proposition 5 shows that there exist two routes r and r' and a vertex v common to both routes such that both sections r_{ov} and r'_{ov} as well as r_{vd} and r'_{vd} are not equal. Since G_1 is series-parallel it attain a well-defined ordering on the vertices. Let A be the first vertex for which the two routes r and r' separate from each other (this vertex can be O itself). Since v is common vertex of theses two routes and $r_{ov} \neq r'_{ov}$ such a vertex exist and its ordering is before vertex v. Because v is a common vertex of r and r' the two routes r and r' have a common vertex A' whose order is after A and before v (it can be v itself). Similarly, we define B as the last vertex for which r and r' coincide (B can be D itself) and B' as the last vertex after v for which r and r' separate (B' can be v itself). The definition of the points A, A', B, B' is illustrated in Figure 6. Next, we show that one of the figures Figure 5a, or 5c, or 5d, or 5e, or 5f, or 5g, or 5h, or 5i is embedded in G. We have the following cases:

- $A = O$, $B = D$, $A' = v$, and $B' = v$: In this case the network shown in Figure 5a is embedded in G. This is because there are two disjoint paths from O to v and from v to D and there is also a path from O to D in G_2. Since any other edge and vertex of the graph belongs to a path that connected O to D, we can construct the graph G by starting from the network shown in Figure 5a and applying the embedding procedure.
- $A = O$, $B = D$, and $A' \neq v$ or $B' \neq v$: In this case the network shown in Figure 5c is embedded in G.
- $A \neq O$, $B = D$, and $A' = v$ and $B' = v$: In this case the network shown in Figure 5d is embedded in G.
- $A = O$, $B \neq D$, and $A' = v$ and $B' = v$: In this case the network shown in Figure 5e is embedded in G.
- $A = O$, $B \neq D$, and $A' \neq v$ or $B' \neq v$: In this case the network shown in Figure 5f is embedded in G.
- $A \neq O$, $B = D$, and $A' \neq v$ or $B' \neq v$: In this case the network shown in Figure 5g is embedded in G.
- $A \neq O$, $B \neq D$, and $A' \neq v$ or $B' \neq v$: In this case the network shown in Figure 5h is embedded in G.
- $A \neq O$, $B \neq D$, and $A' = v$ and $B' = v$: In this case the network shown in Figure 5i is embedded in G.

■

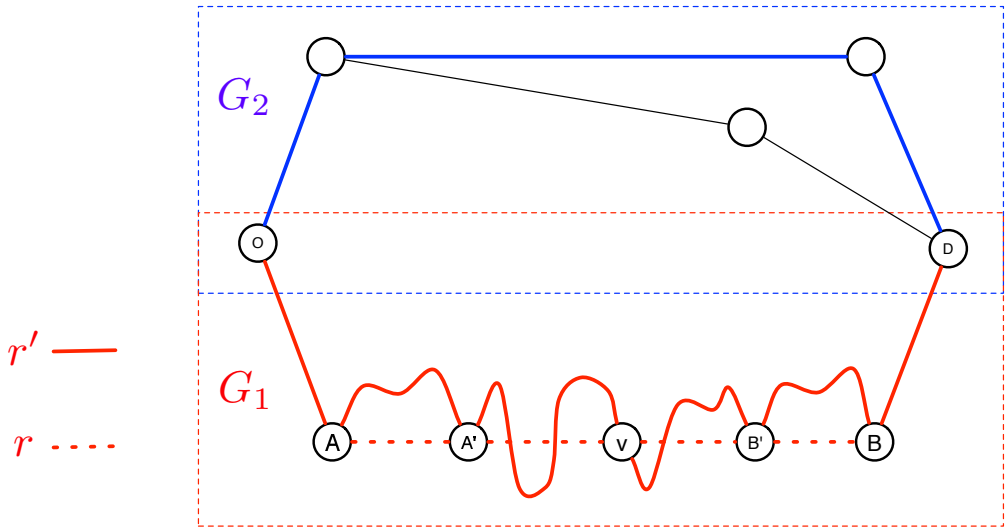

Figure 6: Proof of Theorem 4: G_1 is not LI and G_2 has at least one route from O to D.

Remark 4 Proposition 3 and Proposition 4 can also be stated in terms of a labeled forbidden graph minor, where origin and destination are labeled. A graph H is called a minor of graph G if H can be obtained from G by deleting edges and vertices and contracting edges (merging the two endpoints of the edge)[5]. Here, if we contract an edge that either of its endpoints are O or D, the new vertex will be labeled, O or D. Using this definition, a network with given O and D, is series-parallel if and only if it does not contain the forbidden minor in Figure 4c. Similarly, a network with given O, D pair, is linearly independent if and only if it does not contain the forbidden minors in Figure 4a and Figure 4c. Using labeled forbidden minor definition, a network with given O, D pair is SLI if and only if it does not contain the labeled minors given in Figure 5a and Figure 5b. In this paper, we state our results in terms of graph embedding, to be consistent with the previous related literature on this subject.

Note that the class SLI is a subset of series-parallel graphs and a superset of graphs with linearly independent routes. This is illustrated in Figure 7. The class of SLI graphs plays an important role in our characterization. A natural question is given a graph, how can one determine whether it is SLI or not. We next generalize Proposition 2 and address this question. In particular, we show that in linear time, we can recognize whether a given network is in the class SLI or not.

Proposition 6 *There exists an algorithm that can determine whether a given network G is SLI or not in $O(|\mathcal{E}| + |V|)$.*

Proof Using Proposition 2 we first check whether G is series-parallel or not, which can be done in linear time. We assume G is series-parallel and binary tree decomposition is available (again using Proposition 2, it can be done in linear time). We now start from the root and see if it is S. If it is S, then we have series of two graphs each of them can

[5]When we delete an edge, we keep the endpoints of that edge. When we delete a vertex, we delete all the edges that connect to that vertex. When we contract and edge, we merge the endpoints of it and connect the resulting vertex to all vertices that were connected to either of the two vertices.

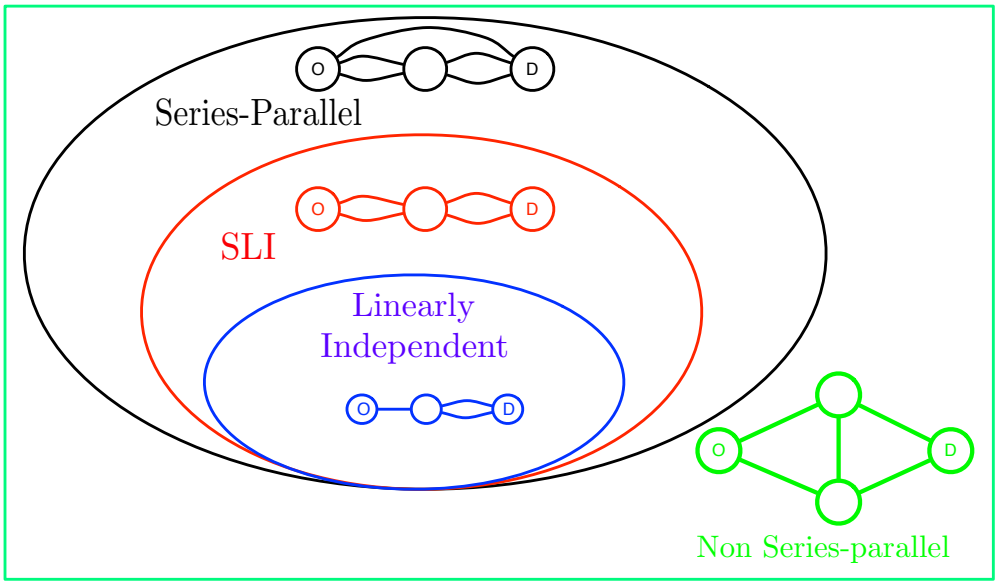

Figure 7: The comparison of class of networks: Linearly Independent $\subset$ Series of Linearly Independent (SLI) $\subset$ Series-Parallel .

be recognized whether they are SLI or not in time linear in their number of edges and vertices and as a result the overall graph can be checked in linear time. Now suppose the root is P. We need to check whether each subtree is LI or not. We next show that given a binary tree decomposition, we can verify whether the underlying graph is LI in $O(V)$. Starting from the root of the tree, if the root has label P, then by induction, for subtree T_1, we can verify whether each subtree is LI in $O(V_{T_1})$. Therefore, in $O(V)$, it can be verified whether the network is LI. If the root is labeled S, then the network being LI is equivalent to have one of the subtrees only labeled S, and the other subtree should be LI. Using any traversing algorithm (breadth first search, or depth first search), one can visit all nodes in the tree, verifying if it only has S labels. Furthermore, by induction we can verify whether each subtree represents an LI network. Therefore, in linear time, one can verify whether the network is LI, completing the proof.

■

5 Informational Braess' Paradox

We first define the classical Braess' Paradox (BP), which is defined for a traffic network with single type of users with $\mathcal{E}_1 = \mathcal{E}$, denoted by $(G, \mathcal{E}_1, s_1, \mathbf{c})$.

Definition 7 (Braess' Paradox (BP)) Consider a traffic network with single information type $(G, \mathcal{E}_1, s_1, \mathbf{c})$. BP occurs if there exists another set of cost functions $\hat{\mathbf{c}}$ with $\hat{c}_e(x) \leq c_e(x)$ for all edges $e \in \mathcal{E}$ and all $x \in \mathbb{R}^+$ such that the equilibrium cost of $(G, \mathcal{E}_1, s_1, \hat{\mathbf{c}})$ is strictly larger than equilibrium cost of $(G, \mathcal{E}_1, s_1, \mathbf{c})$.

BP represents a strongly paradoxical result, because we would ordinarily expect lower cost functions to result in lower equilibrium costs. We next discuss another paradoxical

inefficiency that arises when providing more information to a subset of users in a traffic network with multiple information types.

Definition 8 (Informational Braess' Paradox (IBP)) Consider a traffic network with multiple information types $(G, \mathcal{E}_{1:K}, s_{1:K}, \mathbf{c})$. IBP occurs if there exists an expanded information set $\tilde{\mathcal{E}}_{1:K}$ with $\mathcal{E}_1 \subset \tilde{\mathcal{E}}_1$ and $\tilde{\mathcal{E}}_i = \mathcal{E}_i$, $i = 2, \ldots, K$, such that type 1 equilibrium cost $\tilde{c}^{(1)}$ of $(G, \tilde{\mathcal{E}}_{1:K}, s_{1:K}, \mathbf{c})$ is strictly larger than type 1 equilibrium cost $c^{(1)}$ of $(G, \mathcal{E}_{1:K}, s_{1:K}, \mathbf{c})$.[6]

Similar to BP, IBP represents a strongly paradoxical feature of an equilibrium, because providing more information to a subset of players actually harms these players in equilibrium. It can also be viewed as a generalization of BP, since reducing the cost of an existing edge can be replicated by adding a new edge to a network in which all edges are in the information sets of all users.

The next example illustrates networks in which IBP occurs.

Example 2 In this example we will show that for all graphs shown in Figure 5, there exist an assignment of cost functions along with information sets for which IBP occurs.

(a) Consider the network shown in Figure 5a with the latency functions given by: $c_{e_1}(x) = \frac{1}{2}x$, $c_{e_2}(x) = x + \frac{3}{4}$, $c_{e_3}(x) = \frac{4}{3}x$, $c_{e_4}(x) = 2$ and $c_{e_5}(x) = x$. The information sets are $\mathcal{E}_1 = \{e_1, e_4, e_5\}$, $\mathcal{E}_2 = \{e_2, e_3, e_5\}$, and $\tilde{\mathcal{E}}_2 = \{e_1, e_2, e_3, e_5\}$. For $s_2 = \frac{13}{4}$ and $s_1 = 1$. The equilibrium flows are

$$f^{(1)}_{e_1e_4} = 1, f^{(1)}_{e_5} = 0, f^{(2)}_{e_2e_3} = \frac{3}{4}, f^{(2)}_{e_5} = \frac{10}{4},$$

and

$$\tilde{f}^{(1)}_{e_1e_4} = 0, \tilde{f}^{(1)}_{e_5} = 1, \tilde{f}^{(2)}_{e_2e_3} = 0, \tilde{f}^{(2)}_{e_1e_3} = \frac{6}{4}, \tilde{f}^{(2)}_{e_5} = \frac{7}{4}.$$

The resulting equilibrium costs are $c^{(1)} = c^{(2)} = \frac{10}{4}$ and $\tilde{c}^{(1)} = \tilde{c}^{(2)} = \frac{11}{4}$ and we have $\tilde{c}^{(2)} > c^{(2)}$, and so IBP occurs in this network.

(b) The same setting also shows that IBP occurs in the networks shown in Figure 5c, Figure 5d, Figure 5e, Figure 5f, Figure 5g, Figure 5h, and Figure 5i if we take the cost of extra edges to be zero and include them in all the information sets.

(c) Finally, IBP occurs in Figure 5b. Let $c_{e_1}(x) = c_{e_4}(x) = x + 1$, $c_{e_2}(x) = c_{e_3}(x) = 4 + \frac{1}{2}x$, $c_{e_5}(x) = 1 + x$ and $s_1 = s = 1 - s_2$. Also let $r_1 = e_1e_3$, $r_2 = e_2e_4$, $r_3 = e_1e_5e_4$, $\mathcal{R}_1 = \{r_1, r_2, r_3\}$, $\mathcal{R}_2 = \{r_1, r_2\}$. We define the information sets to be $\mathcal{E}_1 = \{e_1, e_3, e_2, e_4, e_5\}$, $\mathcal{E}_2 = \{e_1, e_3, e_2, e_4\}$, and $\tilde{\mathcal{E}}_1 = \{e_1, e_3, e_2, e_4, e_5\}$. The ICWE for $(G, \mathcal{E}_1, \mathcal{E}_2, s_1, s_2, \mathbf{c})$ is $f^{(1)}_{r_3} = r$ and $f^{(2)}r_2 = f^{(2)}_{r_1} = \frac{1}{2}(1 - s)$ and the ICWE for $(G, \mathcal{E}_1, \tilde{\mathcal{E}}_2, s_1, s_2, \mathbf{c})$ is $f^{(1)}_{r_3} = r$ and $f^{(2)}_{r_3} = (1 - s)$. The corresponding costs are $\tilde{c}^{(1)} = \tilde{c}^{(2)} = 6$ and $c^{(1)} = 4 + 2s$, $c^{(2)} = 5 + s + \frac{3}{4}(1 - s)$. Thus we have $\tilde{c}^{(i)} > c^{(i)}$ for $i = 1, 2$, and IBP occurs in this network as well.

[6] The choice of first type users is without loss of generality, i.e., we assume that the information set of only one type expands and the information set of the rest of the types remain the same.

The network topological conditions that leads to BP has been studied in a seminal work of Milchtaich (2006), in which the author finds sufficient and necessary conditions on network topology under which BP occurs.

Theorem 3 (Milchtaich (2006)) *Consider a traffic network with single information type $(G, \mathcal{E}_1, s_1, \mathbf{c})$. BP does not occur if and only if G is series-parallel. More precisely, we have*

(a) If G is series-parallel, for any assignment of cost functions $\mathbf{c}$, BP does not occur.

(b) If G is not series-parallel, there exists an assignment of cost functions $\mathbf{c}$ in which BP occurs.

We next investigate conditions on the topology of the network under which IBP occurs. Similar to the characterization provided in Theorem 3, we will identify classes of graphs for which IBP does not occur regardless of the cost functions for the edges. Since, as already noted, IBP is a strict generalization of BP, we will see that IBP can occur in a broader class of networks, underscoring the problem mentioned in the Introduction that IBP is likely to be a more pervasive problem.

6 Characterization of Informational Braess' Paradox

In this section, after establishing three key lemmas which underpin the rest of our analysis, we provide our main characterization of IBP, and then extend this result for a more restricted type of change in information sets.

6.1 Three Key Lemmas

The first lemma identifies a property of the traffic network consisting of heterogeneous users over an LI graph.

Lemma 1 *Consider a traffic network with multiple information types $(G, \mathcal{E}_{1:K}, s_{1:K}, \mathbf{c})$, where G is a LI graph. Let $f^{(1:K)}$ and $\tilde{f}^{(1:K)}$ be two (arbitrary) non-identical feasible flows such that $\sum_{i=1}^{K} s_i \geq \sum_{i=1}^{K} \tilde{s}_i$. There exists a route r such that $\sum_{i=1}^{K} f_r^{(i)} > \sum_{i=1}^{K} \tilde{f}_r^{(i)}$ and $f_e \geq \tilde{f}_e$, for all $e \in r$.*

Proof The proof is by induction on the number of edges. We use the characterization of linearly independent networks given in Proposition 4. If G has only one edge, then the result evidently holds. By definition, G is the result of attaching two linearly independent networks (in case of attaching in series one has to be only a single edge, but, for our purpose even attaching two linearly independent networks will work). Let G be the result of attaching G_A and G_B in series. By induction there exist a route r_A in G_A for which $\sum_{i=1}^{K} f_{r_A}^{(i)} > \sum_{i=1}^{K} \tilde{f}_{r_A}^{(i)}$ and $f_e \geq \tilde{f}_e$, for all $e \in r_A$. Similarly, there exist a route r_B in G_B for which $\sum_{i=1}^{K} f_{r_B}^{(i)} > \sum_{i=1}^{K} \tilde{f}_{r_B}^{(i)}$ and $f_e \geq \tilde{f}_e$, for all $e \in r_B$. The route formed by attaching these two routes $r = r_A + r_B$ is a route in G with the desired property. Next, we let G be the result of attaching G_A and G_B in parallel. Since the overall flow in $(G, \mathcal{E}_{1:K}, s_{1:K}, \mathbf{c})$ is larger (not smaller) than the overall in $(G, \mathcal{E}_{1:K}, \tilde{s}_{1:K}, \mathbf{c})$, the overall flow in either G_A or G_B is greater than or equal. Suppose overall flow of G_A is larger. By induction there

exist a route r_A in G_A for which $\sum_{i=1}^K f_{r_A}^{(i)} > \sum_{i=1}^K \tilde{f}_{r_A}^{(i)}$ and $f_e \geq \tilde{f}_e$, for all $e \in r_A$. Note that r_A is also a route in G to complete the proof. ■

The next lemma shows that if G is an LI graph and we expand the information set of one type, e.g., type 1, then the equilibrium cost of at least one of the types does not increase. This result is not sufficient for establishing that IBP does not occur over LI graph because what we need to establish is that it is the equilibrium cost of type 1 that does not increase.

Lemma 2 *Consider a traffic network with multiple information types $(G, \mathcal{E}_{1:K}, s_{1:K}, \mathbf{c})$, where G is a LI graph. Consider an expanded information set $\tilde{\mathcal{E}}_{1:K}$ with $\mathcal{E}_1 \subset \tilde{\mathcal{E}}_1$ and $\tilde{\mathcal{E}}_i = \mathcal{E}_i$, $i = 2, \ldots, K$. Let $c^{(i)}$ and $\tilde{c}^{(i)}$ denote the equilibrium cost of type i users under information sets $\mathcal{E}_{1:K}$ and $\tilde{\mathcal{E}}_{1:K}$, respectively. Then, there exists some $i \in [K]$ such that $\tilde{c}^{(i)} \leq c^{(i)}$.*

Proof Since Lemma 1 holds for any two feasible flows, we can apply it for the equilibrium flows $f^{(1:K)}$ and $\tilde{f}^{(1:K)}$ over the traffic networks $(G, \mathcal{E}_{1:K}, s_{1:K}, \mathbf{c})$ and $(G, \tilde{\mathcal{E}}_{1:K}, s_{1:K}, \mathbf{c})$, respectively (we may view $f^{(1:K)}$ as a feasible flow over the traffic network $(G, \tilde{\mathcal{E}}_{1:K}, s_{1:K}, \mathbf{c})$ as well). It follows that there exists a route r such that $\sum_{i=1}^K f_r^{(i)} > \sum_{i=1}^K \tilde{f}_r^{(i)}$ and $f_e \geq \tilde{f}_e$ for all $e \in r$. From the first inequality it follows that $\sum_{i=1}^K f_r^{(i)} > 0$ which shows at least one of the types, say type i, put a positive traffic on route r. Note that i can be any element of $\{1, \ldots, K\}$ (it can also be 1 as the flows $f^{(1:K)}$ are a feasible flow for the traffic network $(G, \mathcal{E}_{1:K}, s_{1:K}, \mathbf{c})$ as well). We obtain

$$c^{(i)} = c_r \geq \tilde{c}_r \geq \tilde{c}^{(i)},$$

where the first equality follows from $f_r^{(i)} > 0$ and the first inequality follows from $f_e \geq \tilde{f}_e$ for all $e \in r$, and the second inequality follows from the definition of ICWE and the fact that if type i users can use route r in $(G, \mathcal{E}_{1:K}, s_{1:K}, \mathbf{c})$, then they can use it in $(G, \tilde{\mathcal{E}}_{1:K}, s_{1:K}, \mathbf{c})$ as well since the information sets are not smaller in the second game. This completes the proof. ■

The following lemma contains a key result that we will use in the proof of our theorem, Theorem 4. Intuitively, this lemma states that in an LI graph, if we decrease the flow from one part A of the graph and reroute it to another part B of the network (both of which themselves are LI graphs since the original network was LI), then we have that the maximum cost of any route in part A increases relative to the minimum of the cost of any route in part B. This result will enable us to establish that in an LI or SLI network, the reallocation of traffic due to one class of users obtaining more information cannot harm that group.

Lemma 3 *Consider a network with multiple information types $(G, \mathcal{E}_{1:K}, s_{1:K}, \mathbf{c})$, where G is a LI graph. Let $\mathcal{R}_A$ and $\mathcal{R}_B$ denote a partition of routes $\mathcal{R}$ such that $\mathcal{R} = \mathcal{R}_A \cup \mathcal{R}_B$ and $\mathcal{R}_A \cap \mathcal{R}_B = \emptyset$. Let $f^{(1:K)}$ and $\tilde{f}^{(1:K)}$ be two feasible flows. Let $s_A = \sum_{r \in \mathcal{R}_A} \sum_{i=1}^K f_r^{(i)}$, $\tilde{s}_A = \sum_{r \in \mathcal{R}_A} \sum_{i=1}^K \tilde{f}_r^{(i)}$, $s_B = \sum_{r \in \mathcal{R}_B} \sum_{i=1}^K f_r^{(i)}$, and $\tilde{s}_B = \sum_{r \in \mathcal{R}_B} \sum_{i=1}^K \tilde{f}_r^{(i)}$. If $\tilde{s}_A \leq s_A$, and $\tilde{s}_B \geq s_B$, then we have*

$$\max_{r \in \mathcal{R}_A} c_r - \tilde{c}_r \geq \min_{r \in \mathcal{R}_B} c_r - \tilde{c}_r.$$

Before proving this lemma for a general LI graph, let us show it for the special case where G consists of parallel edges from O to D. In this case $\mathcal{R}_A$ and $\mathcal{R}_B$ are two disjoint set of edges from O to D. Since $\tilde{s}_A \leq s_A$, there exist an edge e_A in $\mathcal{R}_A$ such that $\tilde{f}_{e_A} \leq f_{e_A}$. Similarly, since $\tilde{s}_B \geq s_B$, there exist $e_B \in \mathcal{R}_B$ such that $\tilde{f}_{e_B} \geq f_{e_B}$. This leads to

$$\max_{r\in\mathcal{R}_A} c_r - \tilde{c}_r \geq c_{e_A}(f_{e_A}) - c_{e_A}(\tilde{f}_{e_A}) \geq 0 \geq c_{e_B}(f_{e_B}) - c_{e_B}(\tilde{f}_{e_B}) \geq \min_{r\in\mathcal{R}_B} c_r - \tilde{c}_r,$$

which is the desired result.

Proof of Lemma 3 We first note an immediate corollary of Proposition 5:

Claim 1: If a network G is LI then for any vertex such as v, either there exist a unique route from origin to v or there exists a unique route from destination to v. This claim follows from

We now prove the lemma by induction on the number of edges. We have the following cases:

- There exist $r \in \mathcal{R}_A$ such that $c_r \geq \tilde{c}_r$ and $r' \in \mathcal{R}_B$ such that $c_{r'} \leq \tilde{c}_{r'}$. This leads to

  $$\max_{r\in\mathcal{R}_A} c_r - \tilde{c}_r \geq c_r - \tilde{c}_r \geq 0 \geq c'_r - \tilde{c}'_r \geq \min_{r\in\mathcal{R}_B} c_r - \tilde{c}_r,$$

  which concludes the proof in this case.

- For any $r \in \mathcal{R}_A$, we have $c_r < \tilde{c}_r$. Since $\tilde{s}_A \leq s_A$, using Lemma 1 there exists a route $r \in \mathcal{R}_A$ such that the flow on each edge of it from $\tilde{s}_A$ is less than or equal than the flow from s_A. However, we know that the overall cost of $r \in \mathcal{R}_A$ has gone up, i.e., $\tilde{c}_r > c_r$. This implies that the flow from $\tilde{s}_B$ on one edge of r, denoted by e, is more than its flow from s_B. Let $\mathcal{R}_e$ denote the set of routes using edge e. Using Claim 1 there is a unique path from O to the beginning node of e (the unique path is either from O to beginning node of e or from beginning node of e to D; we assume without loss of generality it is the former case). Consider the rest of the route r from v to D, and let $\bar{v}$ denote the closest node to D on r that belongs to some $r' \in \mathcal{R}_e - r$ such that the node before $\bar{v}$ on r' does not belong to r. Note that no other path from $\mathcal{R} - \mathcal{R}_e$ can merge with paths belonging to $\mathcal{R}_e$ between v and $\bar{v}$. Consider the subnetwork between v and $\bar{v}$. Note that this subnetwork consists of only edges belonging to routes in $\mathcal{R}_e$, and let us denote the subparts of the paths in this subnetwork belonging to $\mathcal{R}_A$, and $\mathcal{R}_B$ with $\mathcal{R}'_A$ and $\mathcal{R}'_B$ respectively.

  Using induction on the number of edges, we have

  $$\max_{r\in\mathcal{R}'_A} c_r - \tilde{c}_r \geq \min_{r\in\mathcal{R}'_B} c_r - \tilde{c}_r.$$

  Note that since there is a unique path from O to v, the cost of going from O to v is the same for both flows, and since $\bar{v}$ denotes the last merging point between $\mathcal{R}_e$ and r, all paths in $\mathcal{R}_e$ use the rest of r from $\bar{v}$ to D. Therefore, we have

  $$\begin{aligned}\max_{r\in\mathcal{R}_A} c_r - \tilde{c}_r &\geq (c_{O\to v} - \tilde{c}_{O\to v}) + \max_{r\in\mathcal{R}'_A} c_r - \tilde{c}_r + (c_{\bar{v}\to D} - \tilde{c}_{\bar{v}\to D}) \\ &\geq (c_{O\to v} - \tilde{c}_{O\to v}) + \min_{r\in\mathcal{R}'_B} c_r - \tilde{c}_r + (c_{\bar{v}\to D} - \tilde{c}_{\bar{v}\to D}) \geq \min_{r\in\mathcal{R}_B} c_r - \tilde{c}_r,\end{aligned}$$

  which concludes the proof.

- For any $r \in \mathcal{R}_B$, we have $c_r > \tilde{c}_r$. Since $\tilde{s}_B \geq s_B$, using Lemma 1 there exist a route $r \in \mathcal{R}_B$ such that the flow on each edge of it from $\tilde{s}_B$ is greater than or equal than the flow from s_B. However, we know that the overall cost of $r \in \mathcal{R}_B$ has gone down, i.e., $\tilde{c}_r < c_r$. This implies that the flow from $\tilde{s}_A$ on one edge of r, denoted by e, is less than its flow from s_A. Let $\mathcal{R}_e$ denote the set of routes using edge e. Using Claim 1 there is a unique path from O to the beginning node of e (the unique path is either from O to beginning node of e or from beginning node of e to D; we assume without loss of generality it is the former case). Consider the rest of the route r from v to D, and let $\bar{v}$ denote the closest node to D on r that belongs to some $r' \in \mathcal{R}_e - r$ such that the node before $\bar{v}$ on r' does not belong to r. Note that no other path from $\mathcal{R} - \mathcal{R}_e$ can merge with paths belonging to $\mathcal{R}_e$ between v and $\bar{v}$. Consider the subnetwork between v and $\bar{v}$. Note that this subnetwork consists of only edges belonging to routes in $\mathcal{R}_e$, and let us denote the subparts of the paths in this subnetwork belonging to $\mathcal{R}_A$, and $\mathcal{R}_B$ with $\mathcal{R}'_A$ and $\mathcal{R}'_B$ respectively.

  Using induction on the number of edges, we have

$$\max_{r \in \mathcal{R}'_A} c_r - \tilde{c}_r \geq \min_{r \in \mathcal{R}'_B} c_r - \tilde{c}_r.$$

  Note that since there is a unique path from O to v, the cost of going from O to v is the same for both flows, and since $\bar{v}$ denotes the last merging point between $\mathcal{R}_e$ and r, all paths in $\mathcal{R}_e$ use the rest of r from $\bar{v}$ to D. Therefore, we have

$$\begin{aligned}\max_{r \in \mathcal{R}_A} c_r - \tilde{c}_r &\geq (c_{O \to v} - \tilde{c}_{O \to v}) + \max_{r \in \mathcal{R}'_A} c_r - \tilde{c}_r + (c_{\bar{v} \to D} - \tilde{c}_{\bar{v} \to D}) \\ &\geq (c_{O \to v} - \tilde{c}_{O \to v}) + \min_{r \in \mathcal{R}'_B} c_r - \tilde{c}_r + (c_{\bar{v} \to D} - \tilde{c}_{\bar{v} \to D}) \geq \min_{r \in \mathcal{R}_B} c_r - \tilde{c}_r,\end{aligned}$$

  which concludes the proof.

■

6.2 Characterization of Informational Braess' Paradox

We next present our main result, which states that IBP does not occur if and only if the network is within the class SLI networks. The idea of this result, as already discussed in the Introduction, is that: if in an SLI graph, the reallocation of the flows of users who have obtained new information can never change the allocation of traffic in such a way as to increase costs on all routes in the subgraph to which they are now directing their flows (the "if" part); and an non-SLI graph embeds one of a number of subgraphs (shown in Figure 5), and an IBP can be constructed in each of them (the "only if" part).

Theorem 4 (Characterization of Informational Braess' Paradox) *Consider a traffic network with multiple information sets $(G, \mathcal{E}_{1:K}, s_{1:K}, \mathbf{c})$. IBP does not occur if and only if G is SLI. In other words:*

(a) If G is SLI, for any assignment of cost functions $\mathbf{c}$, IBP does not occur.

(b) If G is not SLI, there exists an assignment of information sets and cost functions $\mathbf{c}$ in which informational Braess' paradox occurs.

Proof of part (a): Without loss of generality assume type 1 is the player with more information, i.e., $\mathcal{E}_1 \subset \tilde{\mathcal{E}}_1$, and for all other types i, $\mathcal{E}_i = \tilde{\mathcal{E}}_i$. To obtain a contradiction, suppose that $\tilde{c}^{(1)} > c^{(1)}$. Note that G is series of several LI graphs and the equilibrium flows on G posses equilibrium flows on each of the linearly independent parts. Since $\tilde{c}^{(1)} > c^{(1)}$, it must be the case that on one of the linearly independent parts we have $\tilde{c}^{(1)} > c^{(1)}$. In the rest of the proof of part (a), we will only consider this linearly independent part.
We partition the set $[K]$ into groups A and B as follows

$$A = \{i \in [k] \ : \ \tilde{c}^{(i)} > c^{(i)}\},$$

and

$$B = \{i \in [k] \ : \ \tilde{c}^{(i)} \leq c^{(i)}\},$$

i.e., set A denotes all types with higher equilibrium cost in the game with higher information, and set B denotes the rest of the types.
We also partition the routes of the network into two subsets $\mathcal{R}_A$ and $\mathcal{R}_B$, where

$$\mathcal{R}_A = \{r \in \mathcal{R} \ : \ \tilde{c}_r > c_r\},$$

and

$$\mathcal{R}_B = \{r \in \mathcal{R} \ : \ \tilde{c}_r \leq c_r\},$$

i.e., $\mathcal{R}_A$ denotes all routes that have higher costs in the game with higher information, and $\mathcal{R}_B$, denote the rest of the routes. We show the following claims:
Claim 1: for any $i \in A$ and $r \in \mathcal{R}_B$, we have $f_r^{(i)} = 0$,
i.e., for a given type, if the equilibrium cost is increased in the game with higher information, then the cost of all routes that the agent was sending flow over is also increased.
This follows since if $r \notin \mathcal{R}_i$, then $f_r^{(i)} = 0$. Otherwise, $r \in \mathcal{R}_i$ which implies $r \in \tilde{\mathcal{R}}_i$

$$c_r \geq \tilde{c}_r \geq \tilde{c}^{(i)} > c^{(i)}$$

where the first inequality follows from the definition of the set $\mathcal{R}_B$ and $r \in \mathcal{R}_B$. The second inequality follows from the definition for ICWE. The third inequality follows from definition of the set A and $i \in A$. The overall inequality shows that $f_r^{(i)} = 0$.
Claim 2: For any $i \in B$ and $r \in \mathcal{R}_A$, we have $\tilde{f}_r^{(i)} = 0$,
i.e., for a given route, if the cost of the route in the equilibrium is increased in the game with higher information, then the equilibrium costs of all types that are using this route in the equilibrium of the higher information game are also increased.
This follows since if $r \notin \tilde{\mathcal{R}}_i$, then $\tilde{f}_r^{(i)} = 0$. Otherwise, $r \in \tilde{\mathcal{R}}_i$ which implies $r \in \mathcal{R}_i$ because $1 \notin B$ and information set of all other types are fixed.

$$\tilde{c}_r > c_r \geq c^{(i)} \geq \tilde{c}^{(i)},$$

where the first inequality follows from the definition of the set $\mathcal{R}_A$ and $r \in \mathcal{R}_A$. The second inequality follows from the definition for ICWE. The third inequality follows from definition of the set B and $i \in B$. The overall inequality shows that $\tilde{f}_r^{(i)} = 0$.
Claim 3: Letting $s_A = \sum_{r \in \mathcal{R}_A} \sum_{i=1}^{K} f_r^{(i)}$, $\tilde{s}_A = \sum_{r \in \mathcal{R}_A} \sum_{i=1}^{K} \tilde{f}_r^{(i)}$, $s_B = \sum_{r \in \mathcal{R}_B} \sum_{i=1}^{K} f_r^{(i)}$, and $\tilde{s}_B = \sum_{r \in \mathcal{R}_B} \sum_{i=1}^{K} \tilde{f}_r^{(i)}$, we have $\tilde{s}_A \leq s_A$ and $\tilde{s}_B \geq s_B$.
This follows from Claims 1 and 2. The flows pushing from $f^{(1:K)}$ over the routes in $\mathcal{R}_A$, i.e., s_A is the entire s_i for all $i \in A$ (Claim 1) and possibly some portion of the flows s_j

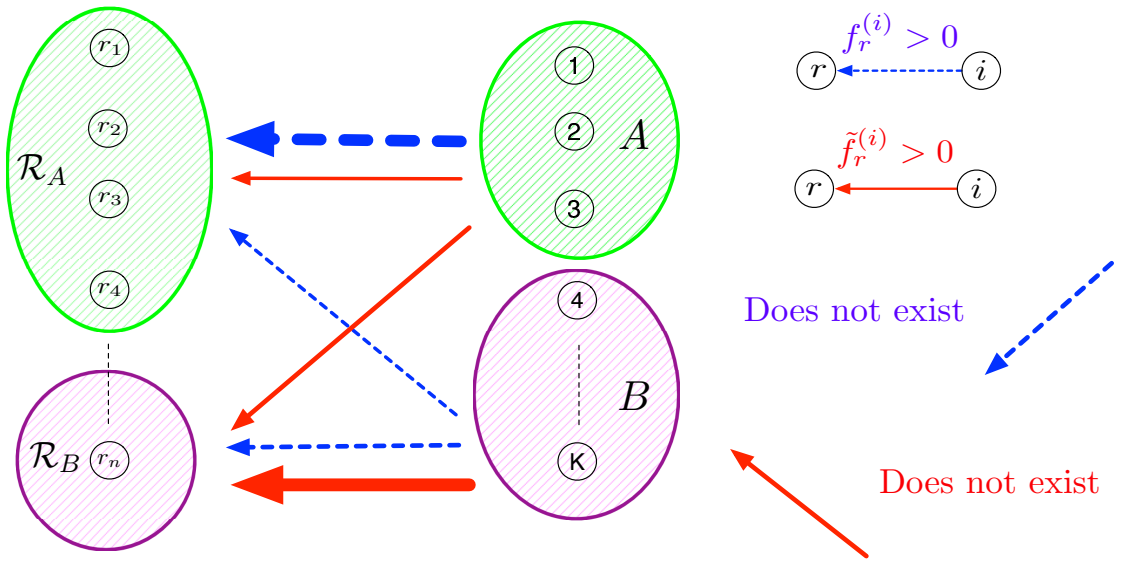

Figure 8: Proof of Theorem 4: the arrows are defined at the right hand side of the figure. The partition of the types into A and B as well as the partition of the routes into $\mathcal{R}_A$ and $\mathcal{R}_B$ as well as the flows are shown in the figure.

for $j \in B$. On the other hand, the flows pushing from $\tilde{f}^{(1:K)}$ over the routes in $\mathcal{R}_A$, i.e., $\tilde{s}_A$ is only some portion of the flows s_i for $i \in A$ (For all $j \in B$ the flow s_j is only pushed through the routes in $\mathcal{R}_B$ (Claim 2)). This shows that $\tilde{s}_A \leq s_A$ which in turn leads to $\tilde{s}_B \geq s_B$ (see Figure 8 for an illustration of the partitioning and the flows).

Note that Lemma 2 shows that there exist type i for which $\tilde{c}^{(i)} \leq c^{(i)}$, which in turn shows that B is nonempty. We also started the proof by the assumption that $1 \in A$, which shows that both A and B are nonempty. Since A is non-empty, because of Claim 1, we have $\mathcal{R}_A \neq \emptyset$ as the flows $f^{(1:K)}$ of the types in A can only go over the routes in $\mathcal{R}_A$. Also, since B is non-empty, because of Claim 2, we have $\mathcal{R}_B \neq \emptyset$ as the flows $\tilde{f}^{(1:K)}$ of the types in B can only go over the routes in $\mathcal{R}_B$.Therefore, we have partitioned the routes of the network into two non-empty sets $\mathcal{R}_A$ and $\mathcal{R}_B$ such that $\tilde{c}_r > c_r$ for all $r \in \mathcal{R}_A$ and $\tilde{c}_r < c_r$ for all $r \in \mathcal{R}_B$. In other words, we have $0 > \max_{r\in\mathcal{R}_A} c_r - \tilde{c}_r$ and $\min_{r\in\mathcal{R}_B} c_r - \tilde{c}_r \geq 0$. We now have all the pieces to use Lemma 3 which yields to

$$0 > \max_{r\in\mathcal{R}_A} c_r - \tilde{c}_r \geq \min_{r\in\mathcal{R}_B} c_r - \tilde{c}_r \geq 0,$$

which is a contradiction. This completes the proof of part a.

Proof of part (b): the proof follows from Theorem 2. First, using Example 2, for all networks shown in Figure 5, there exist an assignment of cost functions and information sets for which IBP occurs. Now, using Theorem 2 one of the figures shown in Figure 5 must be embedded in any network that does not belong to SLI class. We now construct the cost functions and the information sets for the network G following the embedding steps. For a new edge between two nodes put that edge in none of the information sets. For a division of an existing edge assign half of its latency to each of the new edges and update the information set by adding both newly created edge to the same information set as of the original edge. For an extension of origin or destination assign $c(x) = x$ (essentially, any increasing function of x, will work) and update all of the information sets by adding

this edge to them. This assignment will show that the same paradox that was happening over the stating networks (i.e., one of the networks shown in Figure 5) will happen in network G as well. This completes the proof of part b. ■

6.3 IBP with Restricted Information Sets

In this subsection, we show that restricting focus to traffic networks with a much more specific information structure — whereby only one type does not know all the edges and the change in question informs this type of all edges — allows us to establish that IBP does not occur in a larger set of networks. Interestingly, in this case, IBP does not occur in exactly the same set of networks on which BP does not occur, series-parallel networks, though the two concepts continue to be very different even under this more specific information structure. The similarity is that after the change, as in the classic Wardrop Equilibrium setting studied for BP, there is no more heterogeneity among users.

Theorem 5 *Consider a traffic network with multiple information types* $(G, \mathcal{E}_{1:K}, s_{1:K}, \mathbf{c})$. *Suppose that the information sets are as follows:* $\mathcal{E}_i = \mathcal{E}$ *for* $i = 2, \ldots, K$ *and* $\mathcal{E}_1 \subseteq \mathcal{E}$, *and* $\tilde{\mathcal{E}}_1 = \mathcal{E}$. *IBP occurs if and only if the graph is series-parallel. In other words,*

(a) If the graph is series-parallel, then we always have $\tilde{c}^{(1)} \leq c^{(1)}$.

(b) If the graph is not series-parallel, then there exist an assignment of cost functions as well as information sets for which $\tilde{c}^{(1)} > c^{(1)}$.

Proof We fist show two lemmas that we will use in the proof.

Lemma 4 *Consider a traffic network with multiple information* $(G, \mathcal{E}_{1:K}, s_{1:K}, \mathbf{c})$. *If* G *is series-parallel and* $f^{(1:K)} \neq \tilde{f}^{(1:K)}$, *then there exist a route* r *such that* $f_e \geq \tilde{f}_e$ *and* $f_e > 0$ *for all* $e \in r$.

Proof The proof follows by strengthening the claim and then using induction on the number of edges, similar to the proof of Lemma 1. ■

Lemma 5 *Consider a traffic network with multiple information types* $(G, \mathcal{E}_{1:K}, s_{1:K}, \mathbf{c})$ *with equilibrium costs* $c^{(i)}$, $i = 1, \ldots, K$. *Consider a ICWE with flows* $(f^{(1)}, \ldots, f^{(K)})$ *and let* r *be route for which* $f_e > 0$ *for any* $e \in r$. *We have that*

$$c_r \in [\min\{c^{(1)}, \ldots, c^{(K)}\}, \max\{c^{(1)}, \ldots, c^{(K)}\}].$$

Proof We will prove this lemma by induction on the number of edges of G. Since G is series-parallel, then it is either the result of putting two series-parallel networks in series or putting two series-parallel networks in parallel. We consider both cases in the following. If G is the result of putting two series-parallel networks, G_A and G_B in series, then a ICWE for the overall network is obtained by concatenating a ICWE for G_A with a ICWE for G_B (in an arbitrary way). Moreover, we have $c^{(i)} = c_A^{(i)} + c_B^{(i)}$ for $i = 1, \ldots, K$. By induction on the number of edges for the part of r in the graph G_A denoted by r_A, we have

$$c_{r_A} \in [\min\{c_A^{(1)}, \ldots, c_A^{(K)}\}, \max\{c_A^{(1)}, \ldots, c_A^{(K)}\}].$$

We similarly have

$$c_{r_B} \in [\min\{c_B^{(1)}, \ldots, c_B^{(K)}\}, \max\{c_B^{(1)}, \ldots, c_B^{(K)}\}].$$

By adding the previous two relations, we obtain

$$c_r \in [\min\{c^{(1)}, \ldots, c^{(K)}\}, \max\{c^{(1)}, \ldots, c^{(K)}\}].$$

Now suppose that G is the result of putting G_A and G_B in parallel and suppose $r \in G_A$. Let $T = \{i \in [K] \ : \ f_A^{(i)} > 0\}$. For all $i \in T$ since $f_A^{(i)} > 0$, we have $c^{(i)} = c_A^{(i)}$. therefore, by induction we have that

$$c_r \in [\min_{i \in T}\{c^{(i)}\}, \max_{i \in T}\{c^{(i)}\}].$$

Since we have that $\min_{i\in[K]}\{c^{(i)}\} \leq \min_{i\in T}\{c^{(i)}\}$ and $\max_{i\in T}\{c^{(i)} \leq \max_{i\in[K]}\{c^{(i)}$, we obtain

$$c_r \in [\min_{i \in [K]}\{c^{(i)}\}, \max_{i \in [K]}\{c^{(i)}\}],$$

which concludes the proof of lemma. ■

Proof of part (a) of Theorem 5: After expanding type one's information set to $\mathcal{E}$, we obtain $\tilde{c}^{(i)} = \tilde{c}^{(1)}$ for all $i \in [K]$. Using Lemma 4, there exist a route r such that $f_e \geq \tilde{f}_e$ and $f_e > 0$ for any $e \in r$. We have that

$$c_r \geq \tilde{c}_r \geq \tilde{c}^{(i)}, \text{ for } i = 1, \ldots, K,$$

where the first inequality follows form $f_e \geq \tilde{f}_e$, the second inequality follows from the definition of ICWE, and the third inequality follows from $\tilde{\mathcal{E}}_i = \mathcal{E}$ for all $i = 1, \ldots, K$. Since $\mathcal{E}_1 \subseteq \mathcal{E}$, we have $c^{(i)} = c^{(j)} \leq c^{(1)}$ for all $i, j = 1, \ldots, K$. Using Lemma 5, this leads to

$$c_r \in [c^{(i)}, c^{(1)}].$$

Combining the previous two relations, we obtain $\tilde{c}^{(1)} \leq c^{(1)}$.

Proof of part (b) of Theorem 5: In Example 2 we have provided an example showing that IBP can occur over a network which is not series-parallel. The network is shown in Figure 4c.

Suppose that a network G is not series-parallel. Using Proposition 5 G can be constructed from the network shown in Figure 4c by embedding procedure. Using the example we have provided, for this network, there exist an assignment of cost functions and information set for which IBP occurs. We now construct the cost and information set for the network G following the embedding steps. For a new edge between two nodes put the costs of the edge $c(x) = \infty$ and add the edge to both information sets. For a division of an existing edge assign half of its latency to each of the new edges and update the information set by adding both newly created edge to the same information set as of the original edge. For an extension of origin or destination assign $c(x) = x$ (essentially, any increasing function of x, will work) and update all of the information sets by adding this edge to them. This assignment will show that the same paradox that was happening over the stating network (i.e., the network shown in Figure 4c) will happen in network G as well. ■

7 Efficiency of Information Constrained Wardrop Equilibrium

In this section, we provide bounds on the inefficiency of ICWE. The main result is that, though our notion of ICWE is considerably more general than Wardrop Equilibrium as it allows for a rich amount of heterogeneity among users, the worst-case inefficiency remains the same as in standard Wardrop Equilibrium analysis.

We start by defining the social optimum defined as the feasible flow vector that minimizes the total cost over all edges. We consider two different measure of efficiency loss: type-specific efficiency loss defined as the ration of total cost incurred by type i users at the social optimum and ICWE, and aggregate efficiency loss defined as the ratio of total cost experienced by all users at the social optimum and ICWE. We provide tight bounds on both measures of efficiency losses which are realized for different classes of cost functions.

Given a traffic network $(G, \mathcal{E}_{1:K}, s_{1:K}, \mathbf{c})$, we define the social optimum, denoted by $(f_{\text{so}}^{(1)}, \ldots, f_{\text{so}}^{(K)})$, as the optimal solution of the following optimization problem:

$$\begin{aligned} &\min \sum_{e \in \mathcal{E}} f_e \; c_e(f_e), \\ &f_e = \sum_{i=1}^{K} \sum_{\{r \in \mathcal{R}_i \mid e \in r\}} f_r^{(i)}, \\ &\sum_{r \in \mathcal{R}_i} f_r^{(i)} = s_i, \text{ and } f_r^{(i)} \geq 0 \text{ for all } r \in \mathcal{R}_i \text{ and } i. \end{aligned} \tag{7.1}$$

This optimization problem minimizes the total cost over all edges incurred by all users of all types. Under our assumption that each cost function is continuous, it follows that the optimal solution of problem (7.1) and hence a social optimum always exists. We denote the total cost of a feasible flow $f^{(1:K)}$ by

$$C(f^{(1:K)}) \triangleq \sum_{e \in \mathcal{E}} f_e c_e(f_e).$$

Similarly, for a feasible flow $f^{(1:K)}$, we define the total cost incurred by type i users as

$$C^{(i)}(f^{(1:K)}) \triangleq \sum_{e \in \mathcal{E}} f_e^{(i)} c_e(f_e).$$

Consequently, we define the socially optimal cost of type i as $C_{\text{so}}^{(i)} = C^{(i)}(f_{\text{so}}^{(1:K)})$ for $i = 1, \ldots, K$ and the overall cost (over all types) of socially optimum routing as $C_{\text{so}} = C(f_{\text{so}}^{(1:K)})$. Similarly, we define equilibrium cost of type i as $C_{\text{cwe}}^{(i)} = C^{(i)}(f_{\text{cwe}}^{(1:K)})$ for $i = 1, \ldots, K$ and the overall cost (over all types) of ICWE as $C_{\text{cwe}} = C(f_{\text{cwe}}^{(1:K)})$, where $f_{\text{cwe}}^{(1:K)}$ denotes a ICWE.[7]

The following result from Roughgarden and Tardos (2002); Correa et al. (2005) presents bounds on the efficiency loss of Wardrop Equilibrium, which provides bounds on the efficiency loss of ICWE in a traffic network with single information type with $\mathcal{E}_1 = \mathcal{E}$ denoted by $(G, \mathcal{E}_1, s_1, \mathbf{c})$. We use the following result from literature.

[7] Note that $C^{(i)}(f_{\text{cwe}}^{(1:K)})$ is different from equilibrium cost of type i denoted by $c^{(i)}$, as the latter notion is the cost per unit of flow and the former is the aggregate cost. The relation between these two is simply $C_{\text{cwe}}^{(i)} = s_i \; c^{(i)}$ for any $i \in [K]$.

Proposition 7 (Roughgarden and Tardos (2002)) *Consider a traffic network with a single information type $(G, \mathcal{E}_1, s_1, \mathbf{c})$. Let f_{we} be a Wardrop Equilibrium and f_{so} be a social optimum. Then:*

(a) $\inf_{(G,\mathcal{E}_1,s_1,\mathbf{c})\text{: } c_e \text{ convex}} \frac{C_{so}}{C_{we}} = 0.$

(b) Suppose $c_e(x)$ is an affine function for all $e \in \mathcal{E}$. Then, we have $\frac{C_{so}}{C_{we}} \geq \frac{3}{4}$, and this bound is tight.

(c) Let $\mathcal{C}$ be a class of latency functions and let $\beta(\mathcal{C}) = \sup_{c\in\mathcal{C},\ x\geq 0} \beta(c, x)$, where

$$\beta(c, x) = \max_{z\geq 0} \frac{z\ (c(x) - c(z))}{x\ c(x)}.$$

Then we have $\frac{C_{so}}{C_{we}} \geq 1 - \beta(\mathcal{C})$, and the bound is tight.

Our next result shows that Proposition 7 holds exactly for ICWE, indicating that within the class of heterogeneous, information-constrained traffic equilibria we consider, the worst-case scenario occur for networks with homogeneous users.

Proposition 8 *Consider a traffic network with multiple information types $(G, \mathcal{E}_{1:K}, s_{1:K}, \mathbf{c})$. Let f_{cwe} be a ICWE and f_{so} be a social optimum. Then:*

(a) $\inf_{(G,\mathcal{E}_{1:K},s_{1:K},\mathbf{c})\text{: } c_e \text{ convex}} \frac{C_{so}}{C_{cwe}} = 0.$

(b) Suppose $c_e(x)$ is an affine function for all $e \in \mathcal{E}$. Then, we have $\frac{C_{so}}{C_{cwe}} \geq \frac{3}{4}$, and this bound is tight.

(c) Let $\mathcal{C}$ be a class of latency functions and let $\beta(\mathcal{C}) = \sup_{c\in\mathcal{C},x\geq 0} \beta(c, x)$, where

$$\beta(c, x) = \max_{z\geq 0} \frac{z\ (c(x) - c(z))}{x\ c(x)}.$$

Then, we have $\frac{C_{so}}{C_{cwe}} \geq 1 - \beta(\mathcal{C})$, and the bound is tight.

Proof We first show that for any type i, and any feasible flow $f^{(i)}$ for this type, we have

$$\sum_{e\in\mathcal{E}} c_e(f_{e,\text{cwe}})(f_{e,\text{cwe}}{}^{(i)} - f_e^{(i)}) \leq 0. \tag{7.2}$$

The reason is that in ICWE each type uses only the routes with the minimal costs. Therefore, for any type i and any feasible flow $f^{(i)}$ for type i, we have

$$\sum_{r\in\mathcal{R}_i} c_r(f_{\text{cwe}}^{(1:K)}) f_{r,\text{cwe}}^{(i)} \leq \sum_{r\in\mathcal{R}_i} c_r(f_{\text{cwe}}^{(1:K)}) f_r^{(i)}.$$

This leads to

$$\begin{aligned} 0 &\geq \sum_{r\in\mathcal{R}_i} c_r(f_{\text{cwe}}^{(1:K)})(f_{r,\text{cwe}}^{(i)} - f_r^{(i)}) = \sum_{r\in\mathcal{R}_i} \left(\sum_{e:e\in r} c_e(f_{e,\text{cwe}})\right)(f_{r,\text{cwe}}^{(i)} - f_r^{(i)}) \\ &= \sum_{e\in\mathcal{E}} c_e(f_{e,\text{cwe}}) \sum_{r\in\mathcal{R}_i:\ e\in r} (f_{r,\text{cwe}}^{(i)} - f_r^{(i)}) = \sum_{e\in\mathcal{E}} c_e(f_{e,\text{cwe}})(f_{e,\text{cwe}}^{(i)} - f_e^{(i)}), \end{aligned}$$

which is the desired inequality, showing eq. (7.2). We now proceed with the proof.
Part (a) holds because a traffic network with one type is a special case of traffic network with multiple information types and part (a) of Proposition 7 shows the infimum is zero.
Part (b) Using (7.2) for $f^{(i)} = f_{\text{so}}^{(i)}$ for any $i = 1, \dots, K$, and summing over all types $i = 1, \dots, K$, we have

$$\begin{aligned}
C_{\text{cwe}} = \sum_{e\in\mathcal{E}} f_{e,\text{cwe}} c_e(f_{e,\text{cwe}}) &= \sum_{i=1}^{K} \sum_{e\in\mathcal{E}} c_e(f_{e,\text{cwe}}) f_{e,\text{cwe}}^{(i)} \\
&\le \sum_{i=1}^{K} \sum_{e\in\mathcal{E}} c_e(f_{e,\text{cwe}}) f_{e,\text{so}}^{(i)} \\
&\le \sum_{e\in\mathcal{E}} c_e(f_{e,\text{cwe}}) \sum_{i=1}^{K} f_{e,\text{so}}^{(i)} \\
&= \sum_{e\in\mathcal{E}} f_{e,\text{so}} c_e(f_{e,\text{cwe}}) \\
&= \sum_{e\in\mathcal{E}} f_{e,\text{so}} c_e(f_{e,\text{so}}) + \sum_{e\in\mathcal{E}} f_{e,\text{so}} \left(c_e(f_{e,\text{cwe}}) - c_e(f_{e,\text{so}})\right) \\
&\le \sum_{e\in\mathcal{E}} f_{e,\text{so}} c_e(f_{e,\text{so}}) + \frac{1}{4} \sum_{e\in\mathcal{E}} f_{e,\text{cwe}} c_e(f_{e,\text{cwe}}),
\end{aligned}$$

where the last inequality comes from the fact that (assuming $c_e(x) = a_e x + b_e$ for $b_e, a_e \ge 0$)

$$f_{e,\text{so}} \left(c_e(f_{e,\text{cwe}}) - c_e(f_{e,\text{so}})\right) = a_e f_{e,\text{so}}(f_{e,\text{cwe}} - f_{e,\text{so}}) \le \frac{1}{4} {f_{e,\text{cwe}}}^2 a_e \le \frac{1}{4} f_{e,\text{cwe}} c_e(f_{e,\text{cwe}}).$$

The proof of the tightness follows from part (b) of Proposition 7 as traffic network with one type is a special case of traffic network with multiple information types.
Part (c) Using the same argument as in part (b), we have

$$\begin{aligned}
C_{\text{cwe}} = \sum_{e\in\mathcal{E}} f_{e,\text{cwe}} c_e(f_{e,\text{cwe}}) &\le \sum_{e\in\mathcal{E}} f_{e,\text{so}} c_e(f_{e,\text{so}}) + \sum_{e\in\mathcal{E}} f_{e,\text{so}} \left(c_e(f_{e,\text{cwe}}) - c_e(f_{e,\text{so}})\right) \\
&\le \sum_{e\in\mathcal{E}} f_{e,\text{so}} c_e(f_{e,\text{so}}) + \beta(\mathcal{C}) \sum_{e\in\mathcal{E}} f_{e,\text{cwe}} c_e(f_{e,\text{cwe}}),
\end{aligned}$$

where the last inequality comes from the fact that

$$f_{e,\text{so}} \left(c_e(f_{e,\text{cwe}}) - c_e(f_{e,\text{so}})\right) \le \beta(c_e, f_{e,\text{cwe}}) f_{e,\text{cwe}} c_e(f_{e,\text{cwe}}) \le \beta(\mathcal{C}) f_{e,\text{cwe}} c_e(f_{e,\text{cwe}}).$$

The proof of the tightness follows from part (b) of Proposition 7 as traffic network with one type is a special case of traffic network with multiple information types. ■

In concluding this section, we should note that in this environment with heterogeneous users, there are alternatives to our formulation of the social optimum problem, which considers the "utilitarian" social optimum, summing over the costs of all groups. An alternative would be to consider a weighted sum or focus on the class of users suffering the greatest costs. We next illustrate that if we focus on type-specific costs, even with affine cost functions, some groups of users may have worse performance relative to the social optimum than 3/4.

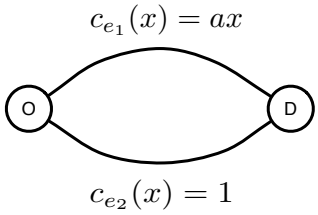

Figure 9: Example 3, type specific price of anarchy versus aggregate price of anarchy.

Example 3 Consider the network shown in Figure 9 with $\mathcal{E}_1 = \{e_1\}$, $\mathcal{E}_2 = \{e_1, e_2\}$. The ICWE is $f^{(1)}_{e_1,cwe} = s_1$ and $f^{(2)}_{e_1,\text{cwe}} = \frac{1}{a} - s_1$, $f^{(2)}_{e_2,\text{cwe}} = s_2 - \frac{1}{a} + s_1$. The equilibrium costs are $C^{(1)}_{\text{cwe}} = s_1$ and $C^{(2)}_{\text{cwe}} = s_2$. The socially optimal flows are $f^{(1)}_{e_1,\text{so}} = s_1$, $f^{(2)}_{e_1,\text{so}} = \frac{1}{2a} - s_1$, and $f^{(2)}_{e_2,\text{so}} = s_2 - \frac{1}{2a} + s_1$. The corresponding costs are $C^{(1)}_{\text{so}} = \frac{s_1}{2}$ and $C^{(2)}_{\text{so}} = \frac{-1}{4a} + s_2 + \frac{s_1}{2}$ (assuming $\frac{1}{2a} \geq s_1$ and $s_2 \geq \frac{1}{a} - s_1$). Therefore, we have

$$\frac{C^{(1)}_{\text{so}}}{C^{(1)}_{\text{cwe}}} = \frac{1}{2}, \ \frac{C^{(2)}_{\text{so}}}{C^{(2)}_{\text{cwe}}} = \frac{\frac{-1}{4a} + s_2 + \frac{s_1}{2}}{s_2}, \text{ and } \frac{C^{(1)}_{\text{so}} + C^{(2)}_{\text{so}}}{C^{(1)}_{\text{cwe}} + C^{(2)}_{\text{cwe}}} = \frac{s_1 + s_2 - \frac{1}{4a}}{s_1 + s_2}.$$

We next show that the ratio of the aggregate costs is greater than or equal to 3/4. We have $s_1 + s_2 \geq \frac{1}{a}$ which leads to

$$\frac{C_{\text{so}}}{C_{\text{cwe}}} = \frac{s_1 + s_2 - \frac{1}{4a}}{s_1 + s_2} = 1 - \frac{1}{4a}\frac{1}{s_1 + s_2} \geq 1 - \frac{1}{4a}a = \frac{3}{4}.$$

However, the type specific ratios can be smaller than $\frac{3}{4}$ as we have that $\frac{C^{(1)}_{\text{so}}}{C^{(1)}_{\text{cwe}}} < \frac{3}{4}$.

8 Concluding Remarks

GPS-based route guidance systems, such as Waze or Google maps, are rapidly spreading among drivers because of their promise of reduced delays as they inform their users about routes that they were not aware of or help them choose dynamically between routes depending on recent levels of congestion. Nevertheless, there is no systematic analysis of the implications for traffic equilibria of additional information provided to subsets of users. In this paper, we systematically studied this question. We first extended the class of standard congestion games used for analysis of traffic equilibria to a setting where users are heterogeneous because of their different information sets about available routes. In particular, each user's information set contains information about a subset of the edges in the entire road network, and drivers can only utilize routes consisting of edges that are in their information sets. We defined the notion of Information Constrained Wardrop Equilibrium (ICWE), an extension of the classic Wardrop Equilibrium notion, and established the existence and essential uniqueness of ICWE.

We then turned to our main focus, which we formulate in the form of Informational Braess' Paradox (IBP). IBP asks the whether users receiving additional information can become worse off. Our main result is a comprehensive answer to this question. We showed that in any network in the series of linearly independent (SLI) class, which is a strict subset of series-parallel network, IBP cannot occur, and in any network that is not in the SLI

class, there exists a configuration of edge-specific cost functions for which IBP will occur. The SLI class is comprised of networks that join linearly independent graphs together, and linearly independent graphs are those for which every pass between origin and destination contains at least one edge that is not in any other such path. This is the property that enables us to prove that IBP cannot occur in any SLI network. We also showed that any network that is not in the SLI class necessarily embeds at least one of a specific set of subgraphs, and then used this property to show that IBP will occur for some cost configurations in any non-SLI graph. We further proved that whether a given graph is SLI can be determined in linear time. Finally, we also established that the worst-case inefficiency performance of ICWE is no worse than the standard Wardrop Equilibrium with one type of users.

There are several natural research directions which are opened up by our study. These include:

- Our analysis focused on the effect of additional information on the set of users receiving the information. For what classes of networks is additional information very harmful for other users? This question is important from the viewpoint of fairness and other social objectives. We may like that users utilizing route guidance systems are experiencing lower delays, but not if this comes at the cost of significantly longer delays for others.

- How "likely" are the cost function configurations that will cause IBP to occur in non-SLI networks. This question is important for determining, ex ante before knowing the exact traffic flows, whether additional information for some sets of users, coming for example from route guidance systems, might be harmful.

- Is there an "optimal information" configuration for users of a traffic network? Specifically, one could consider the following question: given the total flow of K types, $s_1, \ldots, s_K$, find the information sets $\mathcal{E}_1, \ldots, \mathcal{E}_K$ that generate the minimum overall cost for all types in an ICWE. (This question is related to Roughgarden (2001, 2006) who investigate the question of finding the subnetwork of the initial network that leads to optimal equilibrium cost with one type of user).

- Finally, our study poses an obvious empirical question, complementary to similar studies for the Braess' Paradox: are there real-world settings where we can detect IBP?

References

Daron Acemoglu and Asuman Ozdaglar. Competition and efficiency in congested markets. *Mathematics of Operations Research*, 32(1):1–31, 2007.

Daron Acemoglu, Ramesh Johari, and Asuman Ozdaglar. Partially optimal routing. *Selected Areas in Communications, IEEE Journal on*, 25(6):1148–1160, 2007.

Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Eva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.

Richard Arnott and Kenneth Small. The economics of traffic congestion. *American Scientist*, pages 446–455, 1994.

Itai Ashlagi, Dov Monderer, and Moshe Tennenholtz. Two-terminal routing games with unknown active players. *Artificial Intelligence*, 173(15):1441–1455, 2009.

Martin Beckmann, CB McGuire, and Christopher B Winsten. Studies in the economics of transportation. Technical report, 1956.

Priv-Doz Dr D Braess. Über ein paradoxon aus der verkehrsplanung. *Unternehmensforschung*, 12(1):258–268, 1968.

Joel E Cohen and Paul Horowitz. Paradoxical behaviour of mechanical and electrical networks. *Nature*, 352:699–701, 1991.

Joel E Cohen and Clark Jeffries. Congestion resulting from increased capacity in single-server queueing networks. *IEEE/ACM Transactions on Networking (TON)*, 5(2):305–310, 1997.

José R Correa, Andreas S Schulz, and Nicolás E Stier-Moses. Selfish routing in capacitated networks. *Mathematics of Operations Research*, 29(4):961–976, 2004.

José R Correa, Andreas S Schulz, and Nicolás E Stier-Moses. On the inefficiency of equilibria in congestion games. In *Integer Programming and Combinatorial Optimization*, pages 167–181. Springer, 2005.

Stella Dafermos and Anna Nagurney. On some traffic equilibrium theory paradoxes. *Transportation Research Part B: Methodological*, 18(2):101–110, 1984.

Reinhard Diestel. Graduate texts in mathematics: Graph theory, 2000.

Richard J Duffin. Topology of series-parallel networks. *Journal of Mathematical Analysis and Applications*, 10(2):303–318, 1965.

David Eppstein. Parallel recognition of series-parallel graphs. *Information and Computation*, 98(1):41–55, 1992.

Michal Feldman and Ophir Friedler. Convergence to strong equilibrium in network design games.

Michal Feldman and Ophir Friedler. Convergence to strong equilibrium in network design games. *ACM SIGMETRICS Performance Evaluation Review*, 43(3):71–71, 2015.

Marguerite Frank. The braess paradox. *Mathematical Programming*, 20(1):283–302, 1981.

Eric J Friedman. Genericity and congestion control in selfish routing. In *Decision and Control, 2004. CDC. 43rd IEEE Conference on*, volume 5, pages 4667–4672. IEEE, 2004.

Ron Holzman and Nissan Law-yone. Network structure and strong equilibrium in route selection games. *Mathematical Social Sciences*, 46(2):193 – 205, 2003.

Ramesh Johari and John Tsitsiklis. Network resource allocation and a congestion game. In *Mathematics of Operations Research*, 2003.

Frank P Kelly, Aman K Maulloo, and David KH Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research society*, pages 237–252, 1998.

Yannis A Korilis, Aurel A Lazar, and Ariel Orda. Achieving network optima using stackelberg routing strategies. *IEEE/ACM Transactions on Networking (TON)*, 5(1):161–173, 1997.

Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *STACS 99*, pages 404–413. Springer, 1999.

Henry Lin, Tim Roughgarden, and Éva Tardos. A stronger bound on braess's paradox. In *Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 340–341. Society for Industrial and Applied Mathematics, 2004.

Steven H Low and David E Lapsley. Optimization flow controli: basic algorithm and convergence. *IEEE/ACM Transactions on Networking (TON)*, 7(6):861–874, 1999.

Rajiv T Maheswaran and Tamer Başar. Nash equilibrium and decentralized negotiation in auctioning divisible resources. *Group Decision and Negotiation*, 12(5):361–395, 2003.

Marios Mavronicolas, Igal Milchtaich, Burkhard Monien, and Karsten Tiemann. Congestion games with player-specific constants. In *Mathematical Foundations of Computer Science 2007*, pages 633–644. Springer, 2007.

Reshef Meir and David Parkes. Playing the wrong game: Smoothness bounds for congestion games with risk averse agents. *arXiv preprint arXiv:1411.1751*, 2014.

Reshef Meir and David Parkes. Playing the wrong game: Smoothness bounds for congestion games with behavioral biases. *ACM SIGMETRICS Performance Evaluation Review*, 43(3):67–70, 2015.

Igal Milchtaich. Random-player games. *Games and Economic Behavior*, 47(2):353–388, 2004a.

Igal Milchtaich. Social optimality and cooperation in nonatomic congestion games. *Journal of Economic Theory*, 114(1):56–87, 2004b.

Igal Milchtaich. Topological conditions for uniqueness of equilibrium in networks. *Mathematics of Operations Research*, 30(1):225–244, 2005.

Igal Milchtaich. Network topology and the efficiency of equilibrium. *Games and Economic Behavior*, 57:321–346, 2006.

John D Murchland. Braess's paradox of traffic flow. *Transportation Research*, 4(4):391–394, 1970.

GF Newell. Traffic flow on transportation networks. *MIT press series in transportation studies (*, (5), 1980.

E Nikolova and N Stier-Moses. The burden of risk aversion in mean-risk selfish routing. *arXiv preprint arXiv:1411.0059*, 2014.

Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. *Algorithmic game theory*, volume 1. Cambridge University Press Cambridge, 2007.

Ariel Orda, Raphael Rom, and Nahum Shimkin. Competitive routing in multiuser communication networks. *IEEE/ACM Transactions on Networking (ToN)*, 1(5):510–521, 1993.

Michael Patriksson. The traffic assignment problem. 1994.

Georgia Perakis. The price of anarchy when costs are non-separable and asymmetric. In *Integer Programming and Combinatorial Optimization*, pages 46–58. Springer, 2004.

Arthur C Pigou. The economics of welfare. 1920.

J. Riordan and C. Shannon. The number of two-terminal series-parallel networks. *Journal of Mathematical Physics*, (21):83–93, 1942.

Ryan Rogers, Aaron Roth, Jonathan Ullman, and Zhiwei Steven Wu. Inducing approximately optimal flow using truthful mediators. *arXiv preprint arXiv:1502.04019*, 2015.

Robert W Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.

Tim Roughgarden. Designing networks for selfish users is hard. In *Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on*, pages 472–481. IEEE, 2001.

Tim Roughgarden. On the severity of braess's paradox: designing networks for selfish users is hard. *Journal of Computer and System Sciences*, 72(5):922–953, 2006.

Tim Roughgarden and Éva Tardos. How bad is selfish routing? *Journal of the ACM (JACM)*, 49(2):236–259, 2002.

Paul A Samuelson. Spatial price equilibrium and linear programming. *The American economic review*, pages 283–303, 1952.

Sujay Sanghavi and Bruce Hajek. Optimal allocation of a divisible good to strategic buyers. In *Decision and Control, 2004. CDC. 43rd IEEE Conference on*, volume 3, pages 2748–2753. IEEE, 2004.

Lambertus Adrianus Maria Schoenmakers. *A new algorithm for the recognition of series parallel graphs*. Centrum voor Wiskunde en Informatica, 1995.

Yosef Sheffi. Urban transportation networks, 1985.

MJ Smith. The existence, uniqueness and stability of traffic equilibria. *Transportation Research Part B: Methodological*, 13(4):295–304, 1979.

Richard Steinberg and Willard I Zangwill. The prevalence of braess' paradox. *Transportation Science*, 17(3):301–318, 1983.

Jacobo Valdes, Robert E Tarjan, and Eugene L Lawler. The recognition of series parallel digraphs. In *Proceedings of the eleventh annual ACM symposium on Theory of computing*, pages 1–12. ACM, 1979.

Adrian Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *Foundations of Computer Science, 2002. Proceedings. The 43rd Annual IEEE Symposium on*, pages 416–425. IEEE, 2002.

JG Wardrop. Some theoretical aspects of road traffic research. *Proceedings of the Institution of Civil Engineers*, 1:325–378, 1952.

Sichao Yang and Bruce Hajek. Revenue and stability of a mechanism for efficient allocation of a divisible good. *preprint*, 2005.